

2 DVDs Free

LINUX

THE COMPLETE MAGAZINE ON OPEN SOURCE

VOLUME: 06 ISSUE: 09 NOVEMBER 2008 116 PAGES ISSUE# 70

ForYou

Windows Reloaded with Open Source

1 Top 25 Win Apps
25 Open Source Applications
that are a Must for Windows Users

2 Try Linux on Windows
Simple Steps to Install Linux
on Your Windows Itself

3 The First Four Steps
to the World of FOSS

4 Multimedia Simplified
How 'FOSS Multimedia' can
be a Boon for Windows Users

5 A Briefcase for FOSS Apps
Carry Your FOSS Environment
Wherever You Go

6 DVD-2: FOSS on Windows
200-plus Free & Open Source (FOSS)
Applications for Windows Users

Distro DVD
Mandriva 2009




India INR 100
Singapore S\$ 9.5
Malaysia MYR 19



Published by EFY—ISO 9001:2000 Certified



**Feature packed, complete intranet solution for the discerning enterprise/
government organizations/
educational and research institutions.**



**Call for
Free Trial**

- Single SIGN ON system
- GUI based Configuration and administrative program
- Enterprise Back up and Restore functions
- Works with all kinds of Connectivity
- Multi-location support
- Migration from any existing legacy solutions
- Optional Mail/Web hosting services
- 'Optional Clustering Suite' for High availability requirement.

Carizen
SOFTWARE (P) LTD

Chennai : 044-24958222/8228/9296/0610
Mumbai : 022-66492136/37/38
New Delhi : 011 - 26239633/4/5
Bangalore : 080-22208778/9343718998
Hyderabad : 040-27551584/9391370082

Email: info@carizen.com
Web: www.carizen.com



Rainmail Intranet Server

World class network infrastructure at your fingertips

Linux Learning Centre

Pioneers in Training on Linux Technologies

Trained participants from over 36 Countries in 6 Continents

Linux OS Administration & Security Courses for Migration

LLC102: Essentials of Linux OS
LLC103: Linux System & Network Administration
LLC203: Linux Advanced Administration
LLC303: Linux System & Network Monitoring Tools
LLC403: Qmail Server Administration
LLC404: Postfix Server Administration
LLC405: Linux Firewall Solutions
LLC406: OpenLDAP Server Administration
LLC408: Samba Server Administration
LLC409: DNS Administration

Courses for Developers

LLC104: Linux Internals & Programming Essentials
LLC105: Programming with Qt
LLC106: Device Driver Programming on Linux
LLC107: Network Programming on Linux
LLC108: Bash Shell Scripting Essentials
LLC109: CVS on Linux
LLC204: MySQL on Linux
LLC205: Programming with PHP
LLC206: Programming with Perl
LLC207: Programming with Python
LLC208: PostgreSQL on Linux
LLC501: Programming with OpenGL


RHCE Certification Training

RH033: Red hat Linux Essentials
RH133: Red Hat Linux System Administration
RH253: Red Hat Linux Networking & Security Administration
RH300/301: Red Hat Rapid Track Certification Course

RHCSS Certification Training

RHS333: Red Hat Enterprise Security: Network Services
RH423: Red Hat Enterprise Directory Services & Authentication
RHS429: Red Hat Enterprise SELinux Policy Administration

Linux Server Configuration
and Support Services



**Red Hat Fast Track Certification
Course from 10 November 2008**

**RHCE Exams in Bangalore
on 14 & 28 November 2008**

**RH423 from 15 Nov `08
RHS429 from 22 Nov `08
RHS333 from 29 Nov `08**

Linux Support & Solutions
Installation, Setup and Support Solutions
for RedHat, Ubuntu, SUSE, CentOS

For more info log on to:
www.linuxlearningcentre.com
Call: 9845057731 / 9343780054
Email: info@linuxlearningcentre.com

**RED HAT Training Partner
RHCE Examination Centre**



Linux Learning Centre Private Limited

Registered Office & Corporate Training Centre

635, 6th Main Road, (Adj.. Bank of India) Hanumanthnagar, Bangalore 560019

Tel: +91.80.22428538 / 26600839 / 26610999 TelFax: +91.80.26600839

Cell: 9845057731 / 9343780054 Email: info@linuxlearningcentre.com

URL: www.linuxlearningcentre.com

LLC Satellite Centre - Bangalore

1291, 24th Cross, 30th Main,
BSK II Stage, Bangalore 560070

Tel: +91.80.26712928



Windows Reloaded with Free Software

Google Chrome: A Fresh Take On Browsers — 28

Make it Run. Anywhere, Anytime! — 32

Free(dom) Multimedia Powerhouse for Windows — 34

Non-Greasy Ways to Change the Flat Tyres of Your Car — 38

Free Software for Windows: The Top 25 — 40

Run Linux 'in' Windows — 48



FOR YOU & ME

- 18 Of 2009, And All Those Goodies
- 24 Unconferences for the Win
- 26 An Unconference Celebrates the Spirit of Open Source
- 28 Google Chrome: A Fresh Take On Browsers
- 32 Make it Run. Anywhere, Anytime!
- 34 Free(dom) Multimedia Powerhouse for Windows
- 38 Non-Greasy Ways to Change the Flat Tyres of Your Car
- 48 Run Linux 'in' Windows
- 52 Schedule Your Tasks
- 58 Download Them All with Wget

- 60 Parsix GNU/Linux: A Gift From Persia
- 64 The ODF Olympiad: The End of the Opium War?
- 66 Dare To Let That Puppy Loose?
- 72 Are You Ready to Code for Freedom?

GEEKS

- 74 The GRand Unified Bootloader Explained!
- 78 The Building Blocks Of Your Drupal Website—Part 3: The Content Construction Kit
- 88 Programming in Python for Friends and Relations—Python Scripts: Generators Make It Easier to Rescue Systems

CONTENTS

DEVELOPERS

- 84 Linux Threads
- 92 Complex Mathematical Operations with Python

ADMIN

- 96 MaraDNS: A Security-Aware Nameserver
- 98 Go for SOA with Open Source

COLUMNS

- 87 BraveGNUIndia: Building Skills, the Delhi Way
- 91 FreedomYUG: My TV Station
- 102 A Voyage to the Kernel—Day Five: The End of the First Segment
- 106 CodeSport
- 108 The Joy of Programming: What Went Wrong?

REGULAR FEATURES

- 06 Editorial
- 08 Feedback
- 10 Technology News
- 16 Q&A
- 40 DVD Page: Free Software for Windows
- 55 Industry News
- 109 Linux Jobs
- 110 Tips & Tricks

LFY DVD #1



LFY DVD #2



All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under Creative Commons Attribution-Share Alike 3.0 Unported Licence a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-sa/3.0/> for a copy of the licence.

Dear Readers,

Windows on the cover of *LINUX For You*! WHY?

When a member of the *LFY* team suggested this—the rest of us had the same question. But, the idea had merit, as we soon realised.

The primary goal of *LINUX For You* has been to grow the FOSS eco-system. We believe that we have been doing just that by empowering software developers, IT implementers and the community by providing them with information that helps them use FOSS better.

Thanks to the power of the FOSS eco-system, there is a natural increase in the number of FOSS users every year. The new recruits range from young engineers pursuing computer science who are simply mesmerised by the ideology of FOSS, to experienced CIOs who value the freedom and total-cost-of-ownership that FOSS brings to the table.

However, should we remain satisfied with this natural rate of growth? Can't *LFY* do something to accelerate the adoption of FOSS? But that would mean reaching out to 'regular' Windows users and convincing them to start using FOSS. And how could a Linux magazine get itself picked up by Windows users, and influence them to try FOSS software?

"We could do a mega story on FOSS software available on the Windows platform!"

Well, the suggestion started to make sense. There's no doubt that all leading FOSS applications are available on the Windows platform too. All we needed to do was to somehow facilitate Windows users to try out FOSS applications on the platform they were comfortable with—Windows. Once they grow comfortable with them, and appreciate the philosophy behind this eco-system, we believe that many will then try and explore the Linux universe further—and that, in the first place, has always been the key challenge before us.

But, can a Linux magazine talk about Windows software? Good question. But thankfully, the tagline of *LINUX For You*, "The complete magazine on Open Source", gives us that freedom to go beyond the Linux platform. Over the years, we have published quite a few articles on competing platforms like Mac, OpenSolaris, BSD, etc.

So, we got down to business and started preparing for this mega issue. We have two DVDs this time (instead of the usual CD and DVD). One DVD is bundled with over 200 FOSS applications that run on Windows, spread across various categories and sub-categories. The top software from each category has been highlighted in our cover story—to facilitate 'newbies' trying the best of the lot, first. Plus, we have many more articles that are newbie centric, so that this issue can help them get started with their FOSS journey.

So what's in it for the die-hard Linux fans, our regular readers, who're used to getting a new distro to try out every month, that comes bundled on the *LFY* DVD? Good news! It's still there—Mandriva 2009 on the second DVD! And it has OpenOffice.org 3.0 too. :)

We sincerely hope that we have managed to balance our content. Let us know, either way.

Yours truly,



Rahul Chopra

Editor, lfyedit@efyindia.com

PS: For Windows users who are excited enough to try the Linux platform, simply turn to Page 48 to discover simple steps to install Linux 'inside' your Windows system!

Editor

RAHUL CHOPRA

Editorial, Subscriptions & Advertising

DELHI (HQ)

D-87/1, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: (011) 26810602, 26810603
Fax: 26817563
E-mail: info@efyindia.com

BANGALORE

No. 9, 17th Main,
1st Cross, HAL II Stage,
Indiranagar, Bangalore 560008
Ph: (080) 25260023; Fax: 25260394
E-mail: efyblr@efyindia.com

CHENNAI

M. Nackeran
DBS House, 31-A, Cathedral Garden Road
Near Palmgrove Hotel, Chennai 600034
Ph: 044-28275191; Mobile: 09962502404
E-mail: efychn2@efyindia.com

Customer Care

E-MAIL: support@efyindia.com

Back Issues

Kits 'n' Spares

D-88/5, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: (011) 32975879, 26371661-2
E-mail: kits@efyindia.com
Website: www.kitsnspares.com

Advertising

KOLKATA

D.C. Mehra
Ph: (033) 22294788
Telefax: 22650094
E-mail: efycal@efyindia.com
Mobile: 09432422932

MUMBAI

Flory D'Souza
Ph: (022) 24950047, 24928520; Fax: 24954278
E-mail: efymum@efyindia.com

PUNE

Zakir Shaikh
Mobile: 09372407753
E-mail: efypune@efyindia.com

HYDERABAD

P.S. Muralidharan
Ph: 09849962660
E-mail: efyhyd@efyindia.com

Exclusive News-stand Distributor (India)

INDIA BOOK HOUSE PVT LTD

Arch No, 30, below Mahalaxmi Bridge, Mahalaxmi,
Mumbai - 400034 Tel: 24942538, 24925651,
24927383 Fax: 24950392
E-mail: info@ibhworld.com

Printed, published and owned by Ramesh Chopra. Printed at Ratna Offset, C-101, DDA Shed, Okhla Industrial Area, Phase I, New Delhi 110020, on 28th of the previous month, and published from D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020. Copyright © 2008. All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under under Creative Commons Attribution-Share Alike 3.0 Unported License a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-sa/3.0/> for a copy of the licence. Although every effort is made to ensure accuracy, no responsibility whatsoever is taken for any loss due to publishing errors. Articles that cannot be used are returned to the authors if accompanied by a self-addressed and sufficiently stamped envelope. But no responsibility is taken for any loss or delay in returning the material. Disputes, if any, will be settled in a New Delhi court only.

Experience the unique innovative solution

emergic
email2sms

Mail on Mobile

Features

- Receive emails via SMS
- No GPRS required
- No software installation required
- Works with any email service: corporate, POP3 and IMAP accounts.
- Compose, forward & reply immediately
- Receive emails from multiple mailboxes
- Set your own time for receiving emails: DND
- Set filters to receive emails from selected senders
- View attachments without add-on mobile applications



**For Free Lifetime
email2sms service
login to
<http://m3m.in/register4free>**

netCORE
The Innovation Company

sales@netcore.co.in | www.m3m.in

<u>MUMBAI</u>	<u>NEW DELHI</u>	<u>CHENNAI</u>	<u>BANGALORE</u>	<u>HYDERABAD</u>	<u>PUNE</u>	<u>AHMEDABAD</u>	<u>KOLKATA</u>
9820032324	9350101237	9380388685	9341609111	9392224001	9373727934	9824192328	9331420404



You said it...



Just thought that I would thank you for the wonderful article on LUKS (October 2008 issue, Page 58]. I had been contemplating setting up encryption but somehow was not able to find a concise article dedicated to it. Your article surely helped. It's written wonderfully, if only you'd mentioned doing it through a GUI regarding mounting/unmounting and auto mounting. Anyway, thanks to the boost that I got from reading your article, I have configured encryption for my USB disk and memory cards :-). Look forward to more such articles from you.

—Vivek Kapoor, subs.vk@exain.com

Author Sitaram Chamarty

writes: *Thank you for the kind words! I usually avoid using or writing about the GUI for several reasons:*

1. *All recent distros have some sort of GUI for most of these things anyway. I know Fedora 9 does, and Ubuntu 8.04 also, I believe.*
2. *As a result, whatever I write will sound like a user manual for that GUI, plus it will be different for each distro. Since I use only Mandriva, I'd have to spend extra effort learning all the other distros, which takes time.*
3. *It doesn't teach the concepts that lie underneath the GUI—concepts which can carry across distributions and specific desktop environments. Taking encfs as an example, I can extract my files even if I only have access to a back-up copy placed on a server to which I only have command-line access.*

One of the things open source teaches us is not to be dependent on anything more than the bare minimum. A GUI is fine for many

things, but not for backing up, restoring or securing data—these operations are too critical to allow an extra layer of complexity between you and your data. I hope you understand.



First of all, my heartiest congratulations for another successful issue. :-). I was simply amazed by the effort you guys are putting in to deliver a solid magazine. The content, the cover and the disk content were just perfect. Thanks again! I was overjoyed to see seven distros. Well, I wasn't interested in all of them as I have already tested a few like Dream Linux, Mint, gOS, etc. I wasn't interested in OpenGEU, Vector and GoblinX. Though after reading about Pardus, I was tempted to use it, and to my surprise Pardus carried the most innovative things in the LiveCD. I am not talking about just the KDE 4.1.1 desktop, but the overall package and tools shipped. I was stunned by the sheer simplicity of the distro, and the video that introduces Pardus was the icing on the cake. Can you please tell me how to install that LiveCD?

Well, I have a few requests for next month. Being a Mandriva user, can I request you to include its latest offering, Mandriva 2009.0? Apart from that, I am hoping for a detailed review of The GIMP 2.6.1 and Mandriva 2009.0. The new Yellow Pages for FOSS is just a great addition to the magazine. Though I missed the newly started LUG Page in the previous issue, I am pretty satisfied this time. :-). Keep up the good work!

—Shashwat Pant, *by e-mail*

ED: *Your feedback made our day! :-). We had a choice between Slackware or the seven-distros-in-one DVD that we packed. We chose the latter, because that just seemed like more*

fun—and hoped our readers would feel that too. We totally agree with you when you call the Pardus video the 'icing on the cake'; we thought so too. Wonder why other distros don't include such things? As for installing the CD on the computer is concerned, I don't think it's possible; they have a separate installation CD for that.

And, by the time you read this, I'm sure you've found out what our distro DVD contains this month. :-). A review is also included! However, nothing on the new GIMP as yet. But it's coming up!



Hi, I've been a reader of LFY over the last five issues. It's a good magazine, which is resourceful, helpful and has good content. I've been an Ubuntu user since Feisty Fawn. Here in Malaysia, the Linux OS is not popular because it is difficult to configure and it's painful to get the drivers for their peripherals. But I know the ever-resourceful OSS will have a solution for the above-mentioned problems. Here, most Malaysians use 3G for their wireless mobility needs and most of them use Win. XP and Win.Vista, which have a third party driver for PCMCIA cards and Express cards. Could you help me on how to make the 3G cards work on my K/Ubuntu? If possible, could you explain in detail and with photographs? Thank you in advance.

—Mohd Azwar, *Malaysia*

ED: *It's great to hear from a reader in Malaysia! You are lucky to be able to use 3G cards. We in India are yet to see a 3G network come up, thanks to policy delays by the government! It is a challenging task to test something without having the infrastructure, but if we find a solution we'll surely get back to you.*

Please send your comments or suggestions to:

The Editor

LINUX FOR YOU Magazine

D-87/1, Okhla Industrial Area, Phase I,

New Delhi 110020; Phone: 011-

26810601/02/03;

Fax: 26817563; Email: lfyedit@efyindia.com;

Website: www.OpenITis.com



Retain & Retrieve **your mission critical** **emails for several years!!!**

emergic

M@ilArchiva

Email Archiving & Compliance solution

- ✉ Compliance (SOX, HIPPA)
- ✉ Indexing Content & Attachment
- ✉ Insurance against Litigation
- ✉ Encrypted Backup
- ✉ Policy, Security & Retention rules
- ✉ Compression & De-Duplication

Key Benefits

- Seamlessly integrates with any mailing solution (Exchange, Domino, Linux)
- Automated Archiving without end user or administrator intervention
- Quick, Easy Search & Retrieval
- Back up system for lost emails
- No special database required

Partial Clientele List

- Geojit Financial Services Ltd.
- IDBI Capital Ltd.
- Apollo Sindhoori Capital Investment Ltd.
- Flat Products Equipment (I) Pvt. Ltd.
- IndusInd Bank Ltd

netCORE
The Innovation Company

sales@netcore.co.in | www.netcore.co.in/ema

<u>MUMBAI</u>	<u>NEW DELHI</u>	<u>CHENNAI</u>	<u>BANGALORE</u>	<u>HYDERABAD</u>	<u>PUNE</u>	<u>AHMEDABAD</u>	<u>KOLKATA</u>
9820032324	9350101237	9380388685	9341609111	9392224001	9373727934	9824192328	9331420404

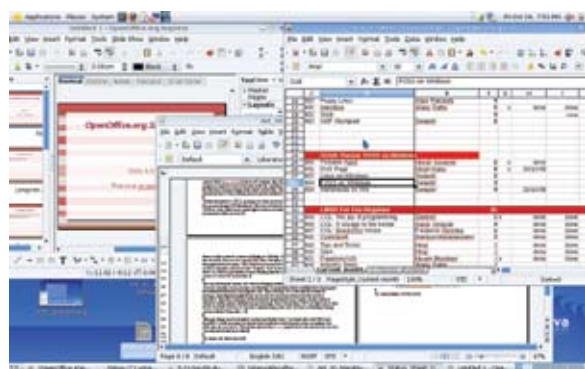
TECHNOLOGY NEWS



OpenOffice.org 3.0 up for grabs

On its eighth birthday, the OpenOffice.org Community announced the release of OpenOffice.org 3.0. The release delivers significant enhancements and advanced, extensible, productivity tools for all users.

OpenOffice.org 3.0 brings in a fresh new look, which includes a new start centre, new splash screen, new icons, and a host of usability improvements. With this release, the basic components, which include a word processor, spreadsheet, presentation, graphics, formula and database capabilities,



can easily be supplemented by extensions downloaded from the OpenOffice.org extensions repository. Instead of feature bloat, OpenOffice.org 3.0 gives users the freedom to configure their suite their way.

However, the basic (core)

components have obviously received a facelift. Writer, among the various other enhancements, has gained a new slider control for zooming which allows multi-page display while editing, improved notes capabilities, and the ability to edit wiki documents for the Web, in addition to conventional office documents. The Calc application brings in an increased capacity of up to 1024 columns per sheet, has gained a powerful new equation solver, and a collaboration feature for multiple users.

Draw now has the ability to cope with poster-size graphics—up to three square metres or 32.3 square feet. Impress has gained a fully-featured table designer inside a presentation. Of course, these features are only the icing on the cake. We recommend you give it a try ASAP. The Mandriva 2009 DVD bundled with this issue of LFY has the OpenOffice.org 3.0 RC2. You can grab the latest version from <http://download.openoffice.org/>

Qt and Qt Jambi updates to v4.4.3

The version numbers of Qt and Qt Jambi have been hiked to 4.4.3 by the project and made available for commercial customers and the open source community. Qt is a cross-platform application framework for desktop and embedded development. Qt Jambi is a version of the Qt for desktop programmers who want to write rich GUI clients using Java. The new versions include updated copyright headers, as well as updated application icons and other graphics to reflect the look and feel of the new Qt brand. These releases do not include bug fixes or optimisations made since the release of version 4.4.2. For detailed information on the fixes and optimisations included, browse the information contained in the developer zone at www.trolltech.com/developer/changes/changes-4.4.3. You can download the open source versions from [trolltech.com/downloads](http://www.trolltech.com/downloads).

Open-Xchange and Funambol to offer push e-mail, PIM functions for handsets

Users of the iPhone, Android, Windows Mobile, BlackBerry and other mobile phones will be able to easily receive push e-mail and synchronise with contacts, calendar and other information on their desktops using open source software from Open-Xchange, a provider of open source collaboration software, and Funambol, a vendor of open source push e-mail and mobile sync.

The Funambol connector for the Open-Xchange Hosting Edition enables push e-mail and the mobile sync of contacts, calendars and tasks with virtually all common wireless phones on the market. Funambol supports more than 1.5 billion devices, including the BlackBerry, iPhone, Windows Mobile, Android, Java ME and SyncML-capable handsets.

Open-Xchange Hosting Edition offers the full set of PIM (personal information manager) functionality—e-mail, calendaring, contacts and task management—together with advanced groupware features such as Documail and Smart Document Sharing. The solution delivers valuable functionality for small businesses and individuals, enabling them to work remotely by providing



low-cost communication and collaboration via the Internet cloud. Leading SaaS providers like 1&1 Internet are using the

joint Open-Xchange Hosting Edition and Funambol push e-mail and sync solution to enable mass market mobile teamwork.

Large **SMP** systems are now as affordable as **HPC**

Problem Statement	Relative Cost		
	SMP	x86 clusters	Tyrone SMP
R&D and manufacturing cost	High	Low	Low
Initial acquisition cost			
Installation and ongoing cost	Low	High	Low
Programming cost			

High Performance Computing (HPC) applications have always needed faster and larger systems, historically served by SMP systems. An alternative for large SMP systems evolved in the late 90s when x86 processors exceeded the performance of RISC processors, promoting the adoption of commodity server clusters connected with high-speed interconnects. Cluster solutions deliver more performance than traditional SMPs, but are more difficult to program, implement and manage.

We at Tyrone systems deliver SMP system, a software based computing-architecture that combines the advantages of shared-memory systems and the price points of commodity clusters by leveraging off-the-shelf x86 components

20-CPU/80-core SMP system with 160GB RAM @ US\$70,000/-*
(20 x Intel Xeon Quad-Core 2.5GHz, 160GB RAM, 20 x 146GB 10000RPM SAS HDDs in 7U rack-mountable)



* Conditions applied

Please do write to us or call us for further details.

Netweb
TECHNOLOGIES
ISO 9002:2000 Company

Corporate Office

B-1/ A25, 2nd Floor
Mohan Co-Operative Industrial Estate,
Mathura Road, New Delhi – 110044

Bangalore : +91 80 41146565/22861407
Chennai : +91 44 43531009/26631975
Kolkatta : +91 33 23232055/23232056


Website : www.netwebindia.com
www.tyronesystems.com
Email : info@netwebindia.com
netweb@vsnl.com

TECHNOLOGY NEWS



The (new and improved) GIMP 2.6

The GIMP team has recently announced the release of a new version of the popular image editing software—version 2.6. In the new version, the toolbox menu bar has been removed and merged with the image window menu bar. To be able to do this, a window called the empty image window has been introduced. It hosts the menu bar and keeps the application instance alive when no images are opened. It also acts as a drag and drop target. When opening the first image, the



empty image window is transformed into a normal image window, and when closing the last image, that window becomes the empty image window.

Important progress towards high bit-depth and non-destructive editing has been made. Most colour operations in GIMP are now ported to the powerful graph based image processing framework GEGL, which means that the internal processing is

being done in 32-bit floating point linear light RGBA. By default, the legacy 8-bit code paths are still used, but a curious user can turn on the use of GEGL for the colour operations with Colours/Use GEGL.

In addition to porting colour operations to GEGL, an experimental GEGL Operation tool has been added, found in the Tools menu. It enables applying GEGL operations to an image and gives on-canvas previews of the results. It is now possible to pan beyond the image border, making image window navigation much less constrained. It is no longer a problem to use the edge of a brush on the edge of an image while being zoomed in, and one can adapt the canvas to any utility window covering parts of the image window.

Get the latest release from www.gimp.org/downloads/

Citrix launches Project Kensho

Citrix Systems has released a technology preview of Project Kensho, its toolkit for the development and deployment of portable virtual machine appliances in enterprises and clouds. Project Kensho aims to offer a multi-hypervisor toolkit that leverages the Distributed Management Task Force (DMTF) Open Virtualisation Format (OVF) to allow ISVs and enterprise IT managers to create hypervisor-independent, portable enterprise application workloads. As a result, virtualised application workloads can be packaged as a secure, portable, pre-configured open standard virtual appliance and be imported and run on Citrix XenServer, Microsoft Windows Server 2008 Hyper-V and VMware ESX virtual environments.

Project Kensho is released as open source software under the Lesser General Public License (LGPL) to accelerate adoption of the industry standard for portable packaging of applications and virtual machines (VMs), as well as management of virtual infrastructure. The Project Kensho technology preview is available for free download on the Citrix Developer Network site at community.citrix.com/display/xs/Kensho.

An all-new Mono 2.0

The Mono project has announced the availability of Mono 2.0, an open source, cross-platform .NET development framework. Mono 2.0 provides all the necessary software to develop and run .NET client and server applications on



Linux, as well as other operating systems. The new Mono 2.0 release is now compatible with the desktop and server components of version

2.0 of the Microsoft .NET framework and features the Mono Migration Analyser (MoMA), an analytical tool for .NET-to-Linux migrations.

More than 2,000 .NET applications are Mono 2.0 compatible with no code changes. Version 2.0 includes performance upgrades—improves scaling and performance for ASP.NET, ADO.NET and Mono runtime; and useful downloads—a virtual machine image that comes with a ready-to-use development environment, as well as many open source Web and desktop .NET applications, including the ASP.NET Starter Kits and other demos. An updated version of the MoMA tool, with improved reporting, is also available and can be downloaded at www.mono-project.com/downloads.

What's new in Python 2.6?

Yes, a new version of one of the most popular programming languages of GenNext, that is Python, was released on October 1. The major theme of version 2.6 has been touted as "... preparing the migration path to Python 3.0, a major redesign of the language." To really understand what's new in Python 2.6, we recommend you visit docs.python.org/whatsnew/2.6.html. Of course, don't forget to grab the new version from www.python.org/ftp/python/2.6/Python-2.6.tar.bz2

**Life as an administrator
is complicated enough...**


**But Backup and Recovery for your
Linux Servers does not need to be.**

NetVault: Backup simplifies backup and recovery
without compromising **Functionality** and **Scalability**

NetVault: Backup provides unmatched Data Protection for all major variants of Linux. We are now offering you the chance to see just how good NetVault is at no cost.

We provide continuous data protection (CDP) for your Linux servers and advanced application protection and recovery for MySQL, PostgreSQL, Oracle, Exchange, DB2 to name just a few.

Permanent FREE Use Edition for Linux is available for download at <http://www.bakbone.com/nvbu/redhat/>



NetVault is a true Linux data protection solution featuring:

- Online backup
- Point and click recovery
- Fully automated protection and recovery support for Linux based applications
- Virtual Tape Library (VTL) with disk staging
- SAN support with LAN free backups.
- Backup duplication for off-site storage of backups
- Support for MySQL native replication, restore to table level and to alternative databases or instances.

BakBone
The Power of Simplicity™

For more information, please contact:

✉ : IndiaSales@bakbone.com
☎ : +91-11-42235156235156

TECHNOLOGY NEWS



Qt Extended (to) 4.4

Qt Software has renamed and launched Qtopia as Qt Extended 4.4, a platform for creating user interfaces and applications for advanced consumer electronics based on Linux. Qt Extended extends the Qt application development framework, bringing increased functionality to embedded Linux. Qt Extended inherits the Qt WebKit integration, which helps device manufacturers enrich applications with live Web content such as online maps, music stores and instant messaging.

Core features of the Qt Extended 4.4 release are: modular architecture for feature selection; an advanced touch-based user interface; IP communications framework based on Telepathy; unified inbox for e-mail, SMS, MMS, IM, etc, plus push IMAP e-mail and other messaging enhancements; and a tool for automated system tests of the target device called Qt UI Test.

Qt Extended 4.4 includes a new Video IP Desk phone reference design that demonstrates an Internet-connected video desk phone UI for finger touch interaction for 3.5" screens. The objective of the reference design is to illustrate the versatility of Qt Extended.



The Qtopia Green phone

Linux development toolkit for ARM

Macraigor Systems, a supplier of embedded debugging tools, and Viosoft Corporation, a provider of integrated embedded Linux software solutions, have jointly announced the availability of an integrated development toolkit for the embedded Linux environment on ARM processors. The toolkit consists of Macraigor hardware debug probes and Viosoft's Arriba Embedded Linux Edition, to enable end-to-end cross-platform development and debugging of Linux applications on ARM.

The toolkit, which is Eclipse-based, includes a full GNU cross-compiler tool-chain for ARM, the Macraigor JTAG probe, the Viosoft VMON2 target-resident debug monitor, and optionally an embedded Linux kernel and file-system for the supported platforms. The toolkit is deployable either as a standalone Eclipse IDE or as a plug-in to Eclipse-based IDEs from third-party vendors. The Macraigor/Viosoft Linux Development Toolkit for ARM is immediately available from Macraigor starting at \$5,995.



CyberLink's digital media solutions go Linux

CyberLink Corp, a solution provider for the connected digital lifestyle, has introduced its high-definition digital media solutions on Linux for netbooks and nettops—CyberLink PowerDVD Linux and PowerCinema Linux. CyberLink PowerDVD Linux is a compact video playback software derived from CyberLink's HD movie player, PowerDVD. To provide the high-definition video and audio experience on Linux PCs, PowerDVD Linux incorporates CyberLink TrueTheater Lighting for automatic video lighting enhancement, and support for CyberLink TrueTheater Surround and Dolby audio technology. The products are, however, only available for OEM licensing.

New JBoss offerings

Red Hat is expanding its enterprise offerings for SOA deployment with the release of JBoss Enterprise SOA Platform 4.3 and JBoss Operations Network 2.1. These new versions will feature cooperative support and will allow for the remote monitoring and management of open source SOA deployments.



JBoss Enterprise SOA Platform 4.3 offers ESB features including gateway listeners, a declarative security model, improved Web services integration and additional scripting languages that will accelerate enterprise adoption of open source SOA by enabling simple Web services integration and deployment. Version 4.3 also features stateful rules services, decision tables and rule agent support that further enable business event processing with an event-driven architecture (EDA), as well as allowing non-developers to construct business rules and enable rapid deployment.

JBoss Operations Network 2.1 provides centralised management for the entire JBoss Enterprise Middleware portfolio, including inventory, administration, deployment and updating of JBoss Enterprise Middleware products and sub-systems. It also features remote platform configuration and deployment, automatic ESB service inventory discovery, monitoring metrics, patch management, and JBoss ESB service monitoring. Built from open source projects such as JBoss ESB, JBoss jBPM and JBoss Rules, the solution offers an enterprise-ready, lightweight footprint and simple installation for low-cost operations. JBoss Enterprise SOA Platform 4.3 and JBoss Operations Network 2.1 are expected to be available by the end of October 2008. For more information visit www.jboss.com.

Microsoft®

BE LOOKED UP TO

Be the one to save the day for developers with a solution, or share something new to become the one everyone would talk about. Start a new project, join an existing one, or download softwares created by fellow developers.

ENTER A NEW HALL OF FAME AT WWW.CODEPLEX.COM,
THE MICROSOFT OPEN SOURCE PROJECT HOSTING WEBSITE.

CodePlex



McCam/MSFT/08



New Horizons KNOW HOW

QNA@NHINDIA.COM



senior corporate
manager—technical
Vipin Sharma
New Horizons India Ltd

Q I have a P4 single computer with 512 MB of RAM. I use CentOS as my OS. After a few minutes of starting my computer, the CPU fan starts to speed up. I have checked all my hardware and found them working fine. One of my friends has suggested that the increase in fan speed is due to the high usage of the CPU. Can you please help me identify the CPU utilisation so that I can troubleshoot?

—Mihir Basu, by e-mail

You are on the right track to troubleshoot. The fan speed is regulated with the CPU temperature. So, as the CPU utilisation increases, the temperature of the CPU also increases and speeds up the fan. To check which process is using up most of your CPU's resources, you can use the top utility. This command provides a dynamic real-time view of a running system. It can display system summary information as well as a list of tasks currently being managed by the kernel. It also gives you information on memory utilisation and the process statistics. Though every distro has a graphical tool for system monitoring, if you are a command line fan, then you have a tool called mpstat that can help you troubleshoot CPU-related performance problems. Other tools that can be used are sar, iostat, vmstat and ps.

Q I work for a company where I need to create a few new users every Monday and keep those valid for a week. Doing this is not a problem, but I want to know if there's a way to force my users to change their login

password on the first login.
—Thomas P., Mangalore

I am sure that you will be familiar with the usermod command. To force users to change their password on the first login, you can use the change command. This changes the number of days between password changes and the date of the last password change. This information is used by the system to determine when users must change their password. To force users to change their password on the first login, you need to run the following after you create the user:

```
chage -d 0 <user_name>
```

...where -d sets the number of days since January 1, 1970 when the password was last changed. The date may also be expressed in the format, YYYY-MM-DD (or what's more commonly used in your area).

Q I use Fedora (GNOME) and a broadband Internet connection. Till a few days back, there was no problem in accessing any site. But, recently, there is a long delay in opening any website. My engineer suggested I try some other DNS, but I am unable to check and change the DNS. Please guide me on how to do this.

—Mini Mathur, by e-mail

Open your terminal and type:

```
cat /etc/resolv.conf.
```

This will give you the IP address of the DNS server your system is using. If you want to change the DNS, then you can open the file in any text editor and change the IP address (do this only as the root).



Q Thanks for the Knoppix Live DVD. It is working well, but the only problem is when using the console in GUI mode, many keys get changed and they mis-type. Please tell me what to do.

—Karthikeyan R., by e-mail

This is possibly because Knoppix is defaulting to a non-US English keyboard layout—could be UK or even German. Although when I tested it, it was defaulting to the US English keyboard on both my systems, it's strange it isn't the same with you. Anyway, simply take a look at the cheat codes and append the US English keyboard option in the boot prompt before you hit 'Enter'. That should take care of the problem you are facing.



New Horizons KNOW HOW



Whenever you need any top-notch IT certification, go for 'New Horizons'!

New Horizons, one of the world's leading IT training companies worth over USD 350 million, has extended its global alliances to India to partner some of the world's leading technology companies. For the corporates, professionals and students, this means that New Horizons' learning experience is further complemented by its faculty, comprising of highly skilled and qualified trainers. As the pioneer of revolutionary learning methods and with a global network in over 60 countries, New Horizons' **Corporate Training, Vendor Certification Programs** and **Career Development Programs** address the needs of corporates, professionals and students to retain and acquire new technology skills that would keep them abreast with the evolving world of IT.

New Horizons has been conferred the 'Best Training Partner, Northern Region' Award by Redhat for four consecutive years.

For any of your top-notch IT certification needs, e-mail at info@nhindia.com And, for corporate enquiries e-mail at corporate@nhindia.com

Vendor Certification Programs

Microsoft
GOLD CERTIFIED
Partners
MCSE, MCSD, .Net

redhat
TRAINING PARTNER
New RHCE 5.0

ORACLE
CERTIFIED PARTNER
DBA, Developer Track

CCNA/CCNP


New Horizons®
CHOOSE. LEARN. SUCCEED

Career Development Programs

N H C N E
New Horizons Certified Networking Engineer
N H C A S
New Horizons Certified Animation Specialist
N H C W S
New Horizons Certified Web Specialist
N H C S E
New Horizons Certified Software Engineer

New Delhi: NBCC Place, Lodhi Road (adjacent to JLN Stadium). Tel.: 24366316, 24365715, 24366349.

India's fastest growing Training & Education brand

www.nhindia.com

*Conditions apply.



Of 2009, And All Those Goodies

Is the newly released Mandriva 2009 good, bad or ugly?

This time round, I switched to Mandriva 2009 on my work laptop as soon as RC2 was out (on September 25, just two days after the official launch of GNOME 2.24). The Mandriva One GNOME Live CD was a 100MB odd lighter than the KDE offering, besides bringing out the latest GNOME offering, so I installed GNOME and then did a:

```
urpmi task-kde4
```

...to get the latest KDE 4.1.1 desktop.

Soon, just before the final release, KDE released an update and the Mandriva packagers didn't take much time to offer v4.1.2 in Cooker (where Mandriva folks cook the distro). Anyone who's sitting with the final release, has this version. Talking about versions, let's have a look at what major version numbers 2009 has in its bag:

- kernel 2.6.27 RC8 (wouldn't 2.6.26 have been a better choice than an RC kernel?)
- KDE 4.1.2 (KDE 3.5.10 is available in the contrib repository for anyone to download)

- GNOME 2.24 (Pidgin is still the default IM client here, unlike Empathy in upstream)
- Firefox 3.0.3 (has been available in other distros for a while, and finally here as well)
- OpenOffice.org 3 (yes, another RC version here; OOo 3 was released after 2009's release)
- Compiz Fusion 0.7.8 (for your 3D desktop needs)

Quite bleeding edge, wouldn't you agree?

As I write this review on my laptop, I am still using the same RC2 base from where I had started—with all the patches and updates installed. The only change I had to make was switch the cooker repository to 2009.0 after the official release on October 9. Things are working quite fine here, so why reinstall?

The third-party PLF repositories (where all the software/patches that can't be put in the official release due to various patents/licence restrictions are hosted) for 2009.0 came live a few days late this time. More about PLF later...

FREE HAS SOMETHING NEW AND BEAUTIFUL TO BEGIN WITH...

...that is, the installer program itself. I've got to say, the Mandriva design team has an eye for beauty, and the installer in just a small example (see Figure 1 on next page). The steps remain pretty much the same compared to the earlier installer. Following is the round-up for the first-timers [I'll do it as I write it, on my home workstation—an assembled AMD Athlon X2 system with 2 GB of RAM, 250 GB and 160 GB hard drives, and the Asus M2N-DVI mother board with that stupid Nvidia graphics card]:

1. Boot screen—Hit *Enter* to get started. This screen also has some other options if we've a very special need.
2. Language—I chose English (British); English (American) is selected by default.
3. Licence agreement—I selected 'Accept' and moved on. After this, it scans the hard disk and pops up the...
4. Install/Upgrade screen—This I presume won't show up if you don't have a previous installation of Mandriva. Mine listed my 2008.1 as something I should upgrade. I clicked on 'Install' for a new install.
5. Keyboard selection—This is funny, and in fact should have followed the Language screen. Anyway, it highlights the UK keyboard as I had chosen British English for language earlier. So, I selected the 'US keyboard'.
6. Partition screen—By default, this has 'Use existing partitions' selected. Other options are 'Use free space', 'Erase and use entire disk' and 'Custom disk partitioning'. This last option suited my needs.
7. Partition screen #2—This lists the hard drive as a graphical rectangle, with colours to indicate the existing partition filesystem types, and white space as free space. My first hard drive has no free space, but my second one (listed as a second tab, labelled *sdb*) has plenty. I selected my exiting home partition and defined mount point as */home* on *sda*, and created a new root partition on *sdb*. It gave a warning as my partition table was about to be written to, with options to say *Cancel* or *OK*. I hit *OK*.
8. Formatting screen—I was asked to chose the partitions I wanted to format. By default it has the root checked to be formatted, and home unchecked. Good enough, so I clicked *Next*. Following which it immediately formatted the partition.
9. Media selection—I was asked if I'd like to copy the entire DVD to my hard disk. I didn't want to. So I clicked *Next*. Following which it downloaded the package data to the memory.
10. Package group selection—This listed all the different repositories available in the DVD, and gave me an option to configure additional media via the CD-ROM or network (http, ftp or nfs). The default setting had none selected, and I clicked *Next*. Following which it looked for available packages.
11. Desktop selection—This has KDE selected by default, plus GNOME and Custom as alternatives. I selected Custom.
12. Package Group selection #2—Too many options here (refer to Figure 1). Defaults seemed okay to me. I only added the GNOME workstation, and the installer listed the total size as more than 4GB. Good enough, I clicked *Next*. Note that there's another option to further fine-tune the installation by selecting the 'Individual package selection' check box.
13. This finally started the installation. (Boy! That's a lot of steps to get here!) It was time to go make myself some coffee and take a break. ;-) *Around 10 minutes later...* Ah! It's done already? That was quick! [Even on my Celeron laptop with 512 MB RAM, it took less than 20 minutes, when I later tested it.]
14. User management—Enter the root password. Enter your user login details. Following which it installs six more packages, and prepares the bootloader.
15. Bootloader installation—Select whether to install it on MBR of *sda*, MBR of *sdb*, first sector of root partition, or skip. The default, which is the first option, suited me, so I moved on. Following which it installs the ACPI package.
16. Summery screen—This is where I can select the time zone, country of residence, bootloader options (again, but with different options), user management (again, if I want to add more users), services to activate/deactivate, keyboard selection (again), mouse, sound card, graphical interface (default has 'Automatic' and looking inside it has the Nvidia GeForce 7050 card selected, which is OK for now as it'll use the free 'vesa' driver. The resolution is Automatic, which I changed to 1280x1024 with 24bpp colours), the network was set to pick up dhcp (fine with me, as I use Airtel broadband), no proxies (I don't have this on my home connection), security level (set to *High* by default, I changed it to *Standard*), and firewall (default is activated; I configured it to allow connections to bittorrent and ssh). That's all!
17. Updates screen—Yes or No. By default it has Yes checked, but I selected No, to get done with it quickly.
18. Congratulations—yes, I'm finally done, and it's time to reboot!

Note: Most of the screens have a help button on the bottom left corner, which brings up a new window explaining every detail of the corresponding screen. Some of the screens also have a Advanced button, to be used to further customise the settings.

Overall, the redesigned installer is a refreshing addition. The only thing left is to reduce the number of steps involved, and organise the procedure a bit more.

One or Free?

Hmm... good question! Well, for the uninformed, One is the LiveCD version of Mandriva, and comes in two flavours: GNOME and KDE4—yes, no KDE3 this time.

This version is a ready solution for desktop users, because of its small size (the included default applications cover almost all our workload) and also because it includes various proprietary drivers, like those for graphics cards



Figure 1: Package group selection screen on the new installer

from Nvidia or ATI, stuff for your wireless chips, etc. Hey, it even includes the Adobe Flash plug-in.

Free is the other version that only includes free software. Another difference between Free and One is that Free comes as a DVD, so expect not just the desktop applications, but a huge pool of server, development, and various other free software that you may or may not need. For that same reason, Free is what's bundled with the November issue of *LFY*. So, let's talk about that version here instead.

However, before I move on, a serious (and rather stupid) issue is that none of the One editions come with printer configuration tools. When it was not included in RC2, I thought it should be fixed in the final release. But, alas, that was only wishful thinking since the final 'Errata' page (which is rather long for a final, and a supposedly stable, release) still listed it as an issue. Anyway, this is, in fact, a serious bug, and should have been taken care of.

While we're on the subject of 'Erratas', if you go through the page, it lists some of the most stupid issues that I can only wonder why they're not fixed before the final release. The Mandriva release team's excuse: the release deadline was up on them, and all of them are just minor issues. My point is: such a long errata page looks plain bad on the project; postponing the release date and fixing at least half of them could have earned them some points. That said, most of the issues listed on Errata will only affect a small sub-set of users, so many don't have to worry. Anyway, let's move on with the installation [see box in previous page] and see how the new version fares.

After installation: It's 2009 in 2008.

The Grub boot loader lists all the previously-installed GNU/Linux versions (and should list Winduhs as well, if you have one). The list of changes (apart from the re-written installer) in 2009, starts from here. On the Grub screen, press *Enter*, and then as the system boots, press *Esc*, to get the verbose mode. Do you notice the difference? Yes, the previous graphical background for



Figure 2: Default KDE desktop and the Mandriva Launcher menu

verbose mode has been replaced by a text-only mode. This is because Mandriva has switched to Splashy [splashy.alioth.debian.org] in 2009 from the bootsplash [www.bootsplash.org] system that was used in earlier releases. The reason cited by the developers is that bootsplash is no more maintained by upstream. That said, Splashy is nowhere as pretty as bootsplash.

The major improvement in this release is the boot speed. Compared to the 2008.1, the difference is quite impressive. I bet even Winduhs never booted this fast even in its glory days. Even on my relatively low-end laptop, it takes around 15-20 seconds in all to load the KDM screen. Anyway, from here on we've to enter a username and password to log in.

The all-new KDE 4.1.2

Let's talk about KDE4 first! Although we've already seen what KDE 4.1 looks like a couple of months back, that distro was a community effort. Mandriva claims their's is the best implementation of KDE4 so far by any distro vendor. Is it? Well, that's also because Mandriva is the first major distribution to officially support KDE 4.1 out-of-the-box. Yes, this is the default desktop (not GNOME), and KDE3 is not even available in the DVD any more, but in their contrib repository online for anyone who wants to settle with that instead. So, is the wait finally over? Is KDE4 usable enough to switch over to?

Mandriva has incorporated its home-grown 'laOra-Qt' style (instead of Oxygen) and the 'Aya' desktop theme. The overall colour settings give you a nice shade of subtle bluish white that is very pleasing to the eye. However, I've got to admit the laOra-Qt style still has some rough edges—I mean, literally, the edges of buttons especially are a bit rough. Or maybe I'm just a big fan of the overall Oxygen style and theme combined, and hence a bit biased. So after changing the style from the System Settings program (a redesigned version of what was known as KControl in KDE3) and also the desktop theme by *right click*→*Desktop Settings* to Oxygen, I was happy.

The desktop, by default, has the the *Desktop* Folder View in the centre of the screen, and the *Trash* applet on the top left corner, where you'd usually expect *Home*

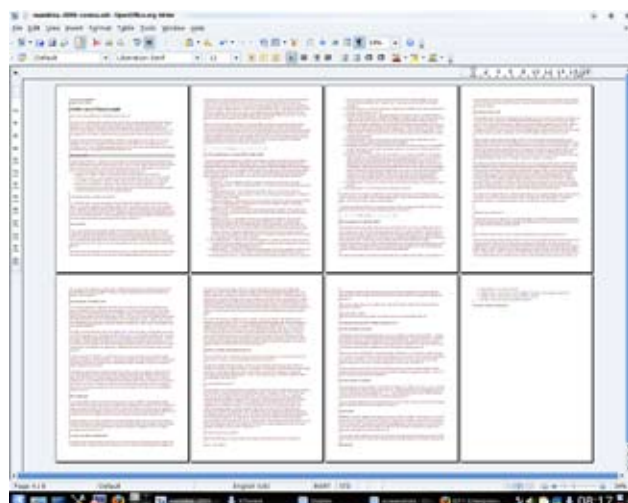


Figure 3: OpenOffice.org 3.0 Writer

(Figure 2). The Mandriva team has put a custom menu called Mandriva Launcher (instead of the Kickoff menu or even the traditional KDE menu) like they always do, and switching to Kickoff is a matter of right clicking and selecting *Switch to Kickoff Menu Style*. The panel has a bunch of icons, which includes the New Device Notifier applet, Show Desktop, System Settings, Mandriva Control Centre, and Firefox. Talking about Firefox, strangely, it has 'Ask' as the default search engine in the search bar, instead of Google. Although, once you switch over to Google, and do a search using it, it'll stay the same. But I still wonder what made them switch to Ask?

The default set of applications cover all areas a typical desktop user would need. I'm sure all of you will first want to check out OOo 3 to begin with. This application's UI has been improved in places, including the addition of the zoom slider in the lower right corner. This is a handy tool as you can have a look at your whole document at one go (Figure 3). Of course, there's much more inside. Go ahead and take a look! One silly complaint: OOo still uses a KDE3 theme, and I would have preferred an Oxygen theme this time.

Like I said earlier, the art team does have an eye for good design and the best example is the default screensaver. It's a slideshow of some of the most beautiful scenery pictures from around the world. I encountered a nice bug here: when the screensaver locks your screen after a while, and you've to enter a password to unlock it again—if you enter a wrong password, it reports some gibberish error message instead of saying 'wrong password'. The first time I encountered it, I got a bit puzzled and thought I'd done something terribly wrong as it asked to me check log files, and what not! You'll encounter this error even on the KDM screen if you enter a wrong password.

The default video player is Dragon, which is quite OK for most purposes, but I'd recommend SMPlayer to any movie buff. SMPlayer being a Qt4-based application, jells quite




Figure 4: GNOME desktop; Nautilus now has tabs

well with the overall KDE4 desktop look and feel, besides using the Mplayer backend, which makes it a much snappier alternative compared to the Mplayer UI or Xine.

The default audio player is a beta version of Amarok2. While this version is almost stable, I couldn't get certain features like Last.fm recommendations to work. It has a lot of cool UI features, but compared to v1.4.x, it still has miles to go.

Overall, KDE4 is a nice experience, and you can surely use it as a KDE3 replacement, now that the KDE PIM suite is also back. In fact, considering KDE4 now even has the spinning cube 3D effect, do you really need Compiz for 3D? That said, I must admit, the shortcut to activate 3D is not a simple *Ctrl+Alt+Left/Right* arrow keys—this works, but first you should hit *Ctrl+F11*. I'd like to see the shortcut keys synced with that of Compiz Fusion—just makes the life of users easier. Note that, 3D and other KDE4 effects are not active by default; activate them from *System Settings*→*Desktop*→*Desktop Effects*.

And, the all-new GNOME 2.24

The GNOME desktop, by default, has the same old 'la Ora' theme as 2008.1 Spring. The default desktop has the regular GNOME icons, and also the other partitions—which is a bad choice as it's unnecessary. But then again, it's the default behaviour of GNOME, I guess. Apart from the theme, the other Mandriva customisations are the Applications menu—you can, of course, revert back to the default GNOME Applications menu by launching MCC and navigating to *System*→*Menu Style*—and an MCC shortcut icon on the panel, next to the default Evolution and Firefox icons. That's all—the rest is all default GNOME.

GNOME 2.24 has introduced some nifty features. One of the very basic, yet important, ones is that the Nautilus file manager now has support for tabs (see Figure 4). As expected, pressing *Ctrl+T* opens a new tab. Although something quite odd about it is how it opens the same location in the new tab—I'd have been happy if it opened the home directory instead.

Another new feature is the compact view. As the name says, it can display more files/directories in the given window without the need for you to scroll because, well,

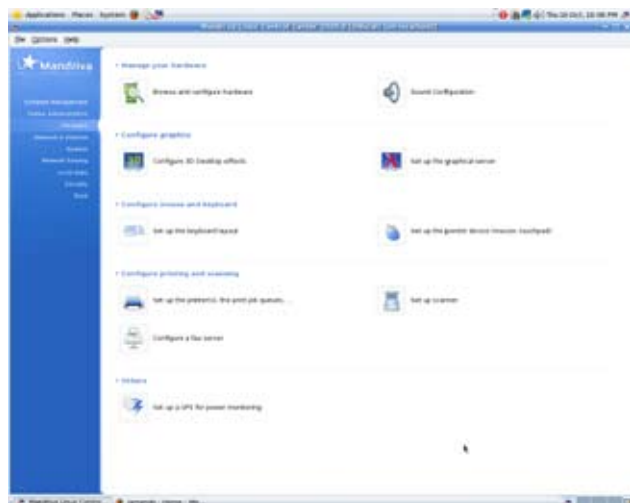


Figure 5: The redesigned Mandriva Control Centre

things are compact here. Although I don't know if it's of any use to me, at the moment.

Another change that GNOME 2.24 introduces is a new instant message client called Empathy. Mandriva, by default, still installs Pidgin. After installing and using Empathy for a while, I've got to say Mandriva has made a smart move by defaulting to Pidgin—Empathy, as of now, is simply too feature-stripped for my tastes.

Overall, there are tiny feature improvements here and there, and over a period of use since the RC2 release, things seemed pretty stable with no application crashes to talk about. However, it seems like GNOME now has become much more resource hungry than what it was a couple of releases back. On my laptop with 512 MB RAM, the desktop didn't seem as responsive as it should have been. I don't know whether the fault lies with GNOME or Mandriva, but I would have liked it a bit snappier. Guess users with 256 MB of RAM will be sort of stuck, eh?

For those with older PCs

Mandriva has added a desktop called LXDE [lxde.org/lxde] to its list of available DEs (desktop environments). This one feels a lot snappier compared to both KDE4 and GNOME. And the best part is, it doesn't look as crippled as the other less-resource hungry DEs make you feel. The only downside is that the desktop is not included in the DVD by default, and you need to install it from their online repositories.

Another alternative is XFCE, packages for which are present in the DVD, for anyone looking for less resource hungry DE alternatives. However, it'd have been better if the Mandriva team included LXDE as well, which scores much higher in terms of usability compared to XFCE. I'd strongly recommend you all try it out.

For those who like to administer

Opening a terminal and editing config files is very cool and has its own place, but it's certainly not for people not




Figure 6: Mandriva update tool

too familiar with the command line, and/or various config file-specific texts. In fact, a lot of us don't even know where a certain config file is located. For such people, the Mandriva Control Centre (MCC) has always been one of the true benefits of using this OS—all system tasks are in one place, there's no need to hunt for them. And what's more? With this release, even *this* has had a facelift (Figure 5).

Although, from past experience you'd believe that anything that ever gets a facelift always packs in a performance penalty as well. That's, however, not the case here—the load time of MCC has reduced to half compared to the earlier editions. As always, all the tools are nicely divided under specific categories based on their functions, and predictably, the names of tools are clear cut to make even the newbies get started with administering their systems using this tool.

One of the best features of Mandriva is how easily you can get started with using those awful Broadcom wireless chips as well. You can simply navigate to the *Network & Internet* category and click the ‘Set up a new network interface (LAN, ISDN, ADSL, ...)’ icon. Once it launches, it will display a set of connection types. Selecting *Wireless* will ask you to enter your DVD, following which it will install a couple of packages, including *ndiswrapper*. In the next screen, select the option that says use a Windows driver, and point it to your Broadcom *inf* file. Browse and locate it, then click *OK*! And there, your laptop’s wireless indicator will come alive. Fill in the ESSID and IP address details to finish the set-up. A simple one-minute procedure, all without fiddling with the command line. Could it be any easier?

For those who like to install and remove

This is the section that MCC defaults to once you launch it. And, seriously, why not? This is the only section that we need to use more frequently than any other. When you first launch it, you are asked whether you'd want

to add the sources from the online repositories. If you have a broadband, hit *Yes*, otherwise *No*. By default, you can only look/search for packages with GUI. Good, if I want to quickly look for GUI apps. Bad if I'm looking for a library. For example, searching for VirtualBox under 'Packages with GUI' option enabled, brings up only one search result. While if you select the 'All' option from the drop-down menu, there are seven packages in all. I needed the VirtualBox guest-additions package too, which I was missing out on if I didn't change the option from *Packages with GUI* to *All*.

Now let's add the online repositories. You already know how to do it! This will add the Main, Contrib and Non Free repositories, and the respective Updates repositories for each. Backports for each are disabled by default, but I'd recommend enabling them. Once done, head over to the Easy URPMI site, and add the PLF repositories as well. From here, you'll get all the media codecs which are either missing, or stripped off of the libraries. Now simply execute the following command to upgrade all packages to the latest updates:

```
urpmi --auto-select
```

While updating your system, this same command also upgrades all the feature-deprived packages to feature-privileged ones by automatically replacing them with packages from PLF repo. And that's it; your Mandriva now supports all sorts of media formats you throw at it. This also gets you the *libfreetype* library from PLF that renders the fonts much better. As far as future updates are concerned, the Mandriva update tool (Figure 6) automatically keeps you covered.

Now that you have the Non Free repository configured, it's time to enable proprietary graphics drivers on the system if you have an Nvidia or ATI chip. Although the steps aren't as obvious as they should be, they are not a geek's play either. Navigate to *MCC*→*Hardware* and click on 'Set up the graphical server' option. Why I said it's not obvious is because, although currently my home system is using the vesa drivers, the 'Graphics Card' section still lists 'NVIDIA GeForce 7050' as the driver. So, there is a possibility of users thinking that the correct driver is installed while the system is still devoid of 3D acceleration.

Anyway, clicking the option, and without changing anything, click *OK* again in the next 'Choose X server' selection screen. This prompts you that there's a proprietary driver available that may support the additional features in your graphics card. Upon clicking 'Yes', it downloads a bunch of packages, including the Nvidia drivers in the case of my home system, from the non-free repository. Once the process of download conclusion and install process is over, you can check for certain options—defaults are good enough. Following this, MCC prompts you to click *OK* to restart X server, or *Cancel* to continue doing your work. Of course, restarting X brings you back all the 3D goodies. :-)

Although, things aren't as bright as they should be

if you have an Nvidia Series 7 card and stick with KDE4—you'll encounter occasional (and sometimes even frequent) screen corruptions, and this really gets on your nerves. The solution given in the *Errata* page didn't solve any issues in my case. You may have better luck with it.

For those who like to serve

Installing the *drakwizard* package brings in two more categories on the side bar of MCC—Sharing and Network Services. While Sharing has wizards to configure an FTP and a Web server, Network Services has wizards to easily configure DHCP, DNS, proxy, time, and SSH servers. There's no need to worry about what packages to install, the wizards take care of installation as well. Hey, there's even a wizard (*drakwizard-trac*) available to set up a Trac project manager system.

What if you need a LAMP stack, not just the Apache Web server? Unfortunately, wizards can't get you everywhere, but this is also as easy as installing a simple meta package like *task-lamp*. Similarly, you can set up a Mandriva Directory Server using *task-mds*, a build system using *task-bs-common*, a cluster using *task-bs-cluster-main*, and more.


For the complete list, launch the software installer (*rpmdrake*), select the 'Meta packages' option from the first drop-down menu on the top left, and then select navigating to *System*→*Servers* in the left-side pane, which lists the categories.

For those who love to hack

As for developers, it includes GCC 4.3.2, Perl 5.10, Python 2.5.2, Ruby 1.8.7, Qt 4.4.3, etc., besides IDEs like Netbeans, Anjuta, and Eclipse. There are a few meta packages for development too, and you can install development packages for C, C++, KDE and Qt, as well as the XFCE build system by navigating to the *Development* section under Meta packages list.

In short, whatever packages you need for your development work is more or less covered, including Web development frameworks like Rails and Django.


Final words

Overall, although this release is a pretty good one, the graphical corruption on Nvidia Series 7 graphics cards, and a pretty lengthy Errata page with minor annoyances are a step back. The experience on the 2008.1 Spring was much, much better. Maybe Spring will bring better news from the French folks.  **END**

REFERENCES

- Home Page: www.mandriva.com
- Release Tour: wiki.mandriva.com/en/2009.0_Tour
- Release Notes: wiki.mandriva.com/en/2009.0_Notes
- Errata: wiki.mandriva.com/en/2009.0_Errata

By: Atanu Datta, LFY Bureau



Unconferences for the Win

A guide to these awesome events.

Let's face it. Conferences are boring. Really! Those massive, speaker-centric events are for the white-collared executives listening to a possibly even more boring person on the stage. Now, this might not hold true for everyone, but it certainly doesn't gel with the current times. Our lives have become faster and time is always short... Which is why the conference needs a makeover.

Un-what?

From Wikipedia, "An unconference is a facilitated, participant-driven conference centred around a theme or purpose." It is not a specific event; rather, it has been applied to a range of gatherings that follow this style. Unconferences are audience-centric events, compared to regular conferences that are speaker-centric events. They are based on the premise that in any professional gathering, the people in the audience—not just those selected to speak on stage—have interesting thoughts, insights, and expertise to share.

The *fundamental* of a conference is: "The sum of the expertise of the people on the stage,

is more than the sum of the expertise of the people in the audience." That's the primary reason why a regular conference is not a discussion; it's more of a monologue. But not an unconference! Everyone who attends an unconference is required to participate in some way—to present, to speak on a panel, to show off a project, or just to ask a lot of questions. Because they do not require the infrastructure and organisation of a full-blown industry gathering, unconferences can happen more frequently. Because the cost to attend is minimal (or non-existent), anyone who wishes to, can come. And because everyone at the unconference participates in some fashion, through interaction, networking and hence participation -- fun and, above all, the exchange of ideas is a given.

How it all began

In his 1997 book, *Open Space Technology*, Harrison Owen discussed many of the techniques now associated with the modern unconference, although his book does not use the term 'unconference'. The term itself was

popularised by Dave Winer, the organiser of BloggerCon.

In his book, Owen gave the following principles that can be applied to unconferences:

- Whoever comes are the right people.
- Whatever happened is the only thing that could have.
- Whenever it starts, is the right time.
- Whenever it's over, it's over.

I know some people will have an issue over the second principle. People say, "No, no... so much more could have happened..." But no, it could not have; at least not under those circumstances.

Owen also gave an accompanying 'Law of Two Feet', which says: "If at any time you find yourself in any situation where you are neither learning nor contributing, use your two feet and move to some place more to your liking."

This is very important to an unconference. People don't need to sit and listen to all the talks. They can move out, gather a group and start impromptu discussions, etc. Such freedom reduces wasting time.

After Owen's work, this method has grown to wild popularity and is used to organise many events around the world. Some examples are BarCamp, BloggerCon, OSScamp, FooCamp, Open Hack Day, etc. The format is immensely popular in the technology community for idea sharing, networking, learning, speaking, demonstrating, and generally interacting with other geeks.

Let's attend one

Since the format is so unconventional, people have difficulty in actually believing that such an event can be feasible and successful. In fact, I remember one of my friends who was to attend an event; he refused to believe that there was no schedule! Unfortunately, you have to participate in such an event to know how wrong you might be! Let's attend one right now...

- Before the event, people usually add their prospective sessions to a website, like a wiki. People might also need to add their name to the attendee list so that it's easier to the organiser in order to keep track of logistics (seating, food, goodies, etc).
- Note: No schedule is decided before hand.
- On event day, people start assembling at the venue. Usually bubbling with activities and anticipation, people get to know each other, have coffee or just hang around.
- Some time later, people get together and decide the schedule of sessions. This process is completely open so that anyone can chip in. Many times, new people add their sessions at that moment itself.
- The format of sessions differs slightly from event to event. Some events have just one room; others have many. Nevertheless, all sessions are very informal and promote participation. People shouldn't feel that they are being taught or lectured!
- Lunch is usually free for all participants. Yaay!
- After the end of the event, all participants usually get together and discuss what else could have been

done and how the event could have been even better. Sometimes, goodies are distributed, and if the event is multi-day, the next day's plan of action is decided.

- Some unconferences include other types of fun events like hackathons, lightning talks and BoF (Birds of a Feather).

I hope you got the gist of the whole process. It is meant to involve everyone. So get up! Learn! Share! After all, not everyone can be an expert on everything.


Challenges involved

Like everything, such events have shortcomings as well. I am not saying these are the negatives of an unconference. These are simply challenges that some people might face. Like the fact that until you attend one, you can't imagine what it's like. Some people think such events are 'directionless'. Others think it's difficult to find the right session, as everything is done on-the-fly. And since most such events are based around a very basic theme, some might seem slightly chaotic.

The Indian scene

India has its fair share of unconferences as well. The BarCamp series around India gets around 300 people per event, and are held in many cities like Delhi, Mumbai, Chennai, etc. These BarCamps focus on varying topics and have a theme, like 'Collaboration', 'Social Media', etc. The participants comprise a wide range of people like business heads, technology geeks, freelancers, students and academia.

Another such series of events is OSScamp. These focus purely on open source software and philosophies. They have been held at Delhi and Mumbai. Bangalore saw the first OSScamp Mobile! which focused on embedded and mobile technologies. OSScamps see programmers, hackers, and FOSS business people from across the country as well. Here, attendance averages around 150.

Unconferences are not just events. They are now being seen as a new form of social organisation. It's the 'we, not I' factor that makes these events what they are. Attend one to experience it yourself!  **END**

REFERENCES:

- *BarCamp website:* barcamp.org
- *OSScamp website:* osscamp.in
- *Unconferences on Wikipedia:* en.wikipedia.org/wiki/Unconference
- *Open Space Technology by Harry Owen:* en.wikipedia.org/wiki/Open_Space_Technology

By: Pratul Kalia. The author is an open source hacker and evangelist. He has been using/tearing up computers since 1996. Currently, he contributes to Drupal, and is a maintainer for Drupal.org and the Ubuntu India forums. He lives on the WWW at <http://pratul.in> and is also known as *lut4rp*.



An Unconference Celebrates the Spirit of Open Source

A report of the OSScamp held in Delhi in September 2008.

The OSScamp community organised its sixth camp on the 27th and 28th of September at the Indian Institute of Technology, Delhi. The community organises such unconferences every few months to propagate open source. This time, the unconference attracted around 130 campers, including those from the LUG at IIT Delhi, other educational institutions and tech professionals. The camp comprised around 50 students, 70 tech professionals and 10 campers from the non-technical *junta*. Around 23 active enthusiasts from outside

Delhi also managed to attend.

OSScamps are community-driven unconferences that are organised on the lines of bar camps and serve as a platform for tech professionals, geeks and geeks-to-be to debate and share their knowledge through demos, presentations and interactions among the campers. It's an endeavour to gather and collaborate on open source technologies and tools.

The two-day event provided a forum for discussions on the impact of new technologies and new ways of thinking about open source phenomenon. The event started with a talk from an organiser of

the camp, Kinshuk Sunil, who spoke on unconferences and their relevance in the technology domain. After this introduction, all the participants were given “One minute of fame” wherein they introduced themselves to fellow campers. A speaker at the camp, Lalit Kumar Shandilya from GeoBeats, shared his views on rapid Web development and what Ruby could do. Lalit also shared his experiences on why GeoBeats is moving from a PHP-based platform to Ruby, and discussed when an organisation could choose Ruby for rapid Web development.

A lunch session was accommodated between the talks, wherein open source was discussed elaborately. This was followed by talks from Supreet Sethi (of MPower Mobile) and Sur Max (of RoR Rockstar) on PostgreSQL and programming in Ruby, respectively. Supreet shared his experiences on how the PostgreSQL database solution is better than MySQL, in the context of Web development, and Sur held an interactive session on programming in Ruby using IRB (Interactive Ruby Shell) for scripting. A talk by an active camper and B.Tech student, Pratul Kalia, marked the end of Day 1. His subject was on how Drupal has powered the new OSScamp website. The talk revolved around the customisation and configuration of Drupal for a community driven event like OSScamps.

Day 2 at the camp commenced with a Linux install fest, where Fedora and Ubuntu were distributed by Sarai and were enthusiastically installed by the campers. Gajendra Khanna, an organiser from IIT, Delhi gave a primer on Python—its definition, features, looping structure, usability criteria and benefits over other scripting languages like Perl, etc. An IIT student, Varun Torka, discussed shell programming and gave a guided demo for writing specific scripts using shell programming, after which Kinshuk gave another talk on open source, covering how open source had evolved from being a development methodology to being a culture now.


Cross-platform and cross-language programming was discussed by Gora Mohanty, a programmer for Sarai's Ind-Linux project. He specifically talked about a project Sarai has done for online Hindi spell-checking using Mono, Aspell and Swig. Gora also gave a demo on how Swig could be used to translate program code written in C++ into C# and Python. Post lunch, open source business models were discussed by Pradyot Sahu, director, Apxpertise, who elaborated on examples of those in the process of establishing open source businesses, the associated concerns and suggestions. An enriching session was a demo of *Big Buck Bunny*, a short movie made by the Blender Foundation and the Blender Community to showcase the power of the Blender engine. The movie was played to the enthusiastic applause of the campers and was followed by

talks from Vikas Yadav and Pradyot.

Vikas gave a demo of OSM (Open Street Map), using GPS devices to track new locations and also shared how OSM is completely open source in comparison with Google Maps, Yahoo Maps and MapMyIndia API. Pradyot talked on open innovation and how it contributes to humanity's knowledge-wealth. The last session of Day 2 had two active Drupalers, Pratul and Gurpratap Singh (from Hoshiarpur), going through a demonstration module of development for Drupal and discussing Drupal tools. Expertise, enthusiasm and interest amongst campers accentuated sharing and learning.

The OSScamp community wants to thank the sponsors for their support. A warm thanks to OSSCube Solutions for the generous financial support, Directi for the banners and stickers, and ZyXEL for extending the network support.

With the successful conclusion of OSScamp Delhi, September 2008, the community is now planning to organise a camp aligned with the IIT technology fest, Tryst, next year. The event would play host to almost 100 technical institutes in India plus some of the most influential technology companies in the world. At the camp, the inception of OSScamp Nainital Community also took place, with student volunteers taking on the responsibility of forging an active FOSS-oriented community around the Nainital area.

For more information on the event, visit the website at www.osscomp.in 

By: Priyanka Jain





Google Chrome™

A Fresh Take On Browsers

The release of Google Chrome in early September heralds a fresh take on browsers. What's more, Chrome is an end-to-end open source browser with the complete code base available as the Chromium Project.

Google Chrome is now available as a beta download for MS Windows XP and Vista only. Install it and you are ready to begin your journey into the latest offering from Google.

Chrome has incorporated a lot of user interface changes compared to other browsers. The first impression is that it has a minimalist user interface like many of its Google stable mates. You will not, for example, see the menu bar or the tool bar.

Let's now embark upon a brief look at some of the compelling usability features of Chrome.

What's in it

Although details on all the features of Google Chrome are listed at www.google.com/chrome/intl/en/features.html, let me introduce you to some of the most promising ones.

Omnibox: Chrome combines the address bar and search box into an 'omnibox'. As you type in the box, Chrome uses your browsing

history and Google Suggest [labs.google.com/suggestfaq.html] to incrementally display suggestions on what the intended destination is likely to be (Figure 1). If you type a domain name, Chrome takes you there. If you type a phrase of text, Chrome uses the default search engine to look up the phrase.

Application shortcuts: You can create an application shortcut by clicking on the first button on the tab of the opened Web application (Figure 2). Using this feature you can place a desktop/start menu shortcut to a Web application of your choice. Whenever you launch the Web app using the shortcut, you will not see any browser features like the *Back/Forward/Reload* buttons. All you see is the Web application. In a sense the website starts behaving like a true Web application on the desktop and does not give the impression of running on a browser. I like this feature since I often have applications like GMail running throughout the day and all I need is the Web application and not the browser with all its features. Chrome uses Google Gears

[gears.google.com] to enable this functionality.

Crash control: If one out of the 10 open tabs crashes, it won't take down the Web browser as a whole—that is, the other nine tabs are usable and safe.

Dynamic tabs: Google Chrome has truly dynamic tabs enabling a seamless tab-based browsing experience.

Incognito mode: This is stealth mode browsing. What you do in this mode is not recorded in the browsing history.

No tie-ups to Google services: Google Chrome is not tightly tied in to any of the other Google services, such as Google Search. You can easily change or add a new default search engine from the Google Chrome Options under the *Basics* tab. Figure 3 shows the *Search Engines* selection window.

Task manager: Multiple tabs can often place high demands on your system. The task manager (*Control the Current Page*→*Developer*→*Task Manager*) shows details about the currently open Google Chrome tabs and plug-ins.

New tab window: When you open a new tab, it is because you want to go somewhere—maybe a page you visited earlier. Google Chrome understands this and presents you with a boxed view of your 'most visited' history and recent bookmarks (Figure 4).

No Auto Feed Discovery: This feature is not present. We don't know why, especially given that Google has feed services like Google Reader and FeedBurner.

A sophisticated core

Google Chrome combines a sophisticated core with a minimalist user interface. We shall now take a look at some of its core architectural features:

Webkit rendering engine: Chrome uses the open source browser engine Webkit [webkit.org], which has proven to be sufficiently fast on desktops and mobile phones alike. This was the main reason why Chrome chose the engine. Using Webkit also saves Web developers from worrying about yet another rendering engine for their content to look uniform. Webkit is already used by Apple Safari [www.apple.com/safari].

Multi-process architecture: With its multi-process architecture, Chrome ensures separate operating system processes for different Web applications. Each Web application has its own space. During start-up there is a browser process that forks a new renderer process every time a new tab is opened; or a new plug-in process every time a plug-in, either Java or Flash, is opened.

For example, Figure 5 shows the process tree when I had four tabs—each running a different Web application, with one of them even running the Flash plug-in playing a movie.

The visual in Figure 5 explains the multi-process hierarchy. At the top level there is the parent process, which is called the browser process and then we have the child processes composed of renderer and plug-in



Figure 1: The omnibox



Figure 2: Create an application shortcut

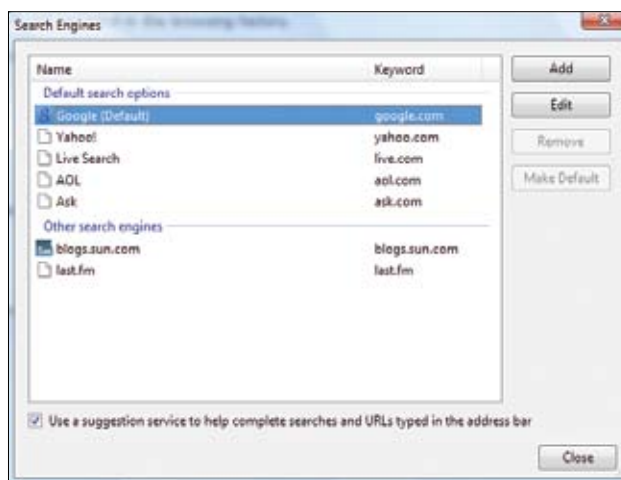


Figure 3: Select your default search engine, or add new ones



Figure 4: A new tab presents a list of most visited websites

processes. Please refer to the blog post titled “Multi-process Architecture” at blog.chromium.org/2008/09/multi-process-architecture.html for a detailed discussion on the topic. All the currently running processes can be viewed using the task manager, as described earlier.

Sandboxed security architecture: Chrome uses a modified version of the BIBA security model [en.wikipedia.org/wiki/Biba_Model]. In Chrome, the browser process and the renderer process have different access privileges. Rendering engines live in a sandbox and are



chrome.exe (3128)	Google Chrome
chrome.exe (1220)	Google Chrome
chrome.exe (2764)	Google Chrome
chrome.exe (1028)	Google Chrome
chrome.exe (1428)	Google Chrome
chrome.exe (4072)	Google Chrome

Figure 5: The process tree of different tabs

under the supervision of the browser process. A detailed report on the security architecture of Chrome is at crypto.stanford.edu/websec/chromium.

V8—JavaScript Virtual Machine: JavaScript is the major driving force of the evolving Web. Realising this, the Chrome team wrote a new JavaScript engine called V8 [code.google.com/apis/v8] to give superior JavaScript performance. V8 is a major architectural characteristic that contributes to its speed. V8 is an open source project, which can be used by developers to write better-performing JavaScript-based Web applications.

DNS pre-fetching or pre-resolving: To speed up browsing, Chrome resolves domain names before the user navigates, typically while the user is viewing a Web page. As the user scrolls through the page containing multiple unvisited domain names, Google uses the computer's DNS resolution mechanism to resolve them and store them

in its 'cache', so that when the user clicks on them, the transfer to that page is usually fast, as the DNS resolution has already been done.

For more details on the architectural features, visit blog.chromium.org.

With Chrome, Google asks users to download a new Web Browser. There is a new interface and new technological advances in Chrome. At this point of time, the browser has features that are ideal for taking it out on a test drive.

Issues

Google Chrome is a beta product, so please do not expect an all-smooth performance. Some known issues are listed at www.google.com/support/chrome/bin/request.py?contact_type=known_issues_2.

Further Information

- Google Chrome: <http://www.google.com/chrome>
- Google Chrome Comic Book: <http://www.google.com/googlebooks/chrome/index.html>



By: Amit Kumar Saha is passionate about writing and blogging. He works for Sun Microsystems, India. All views/opinions expressed in this article are his own and not of his employers. He blogs at <http://blogs.sun.com/amitsaha>, and can be reached at amitsaha.in@gmail.com

LINUX

ForYou

Calling...

Share your knowledge

Inviting FOSS experts to write articles on their area of interest

LFY covers a myriad of topics—network management, software development, embedded systems, community issues, and even hands-on guide for newbies. If you've got an interesting topic, let us know. Thanks to the launch of www.openITis.com (aka linuxforu.com) we are now trying to extend our content portfolio related to Linux & Open Source.

To know more on how to become an LFY author, contact us at

lfyedit@efyindia.com

Few topics that top our list:

- ◆ Tips 'n' Tricks for software developets or IT implementers
- ◆ Cool tweaks for FOSS enthusiasts
- ◆ FOSS on mobile
- ◆ Virtualisation (Implementation)
- ◆ OpenJDK or Java on Linux
- ◆ OpenSolaris (software development)
- ◆ How can I do 'that' on Linux
- ◆ Reviews of latest open source projects & tools

**Server Consolidation,
Containment,
with VMware Virtual Infrastructure**

**Business Continuity,
And Disaster Recovery,**

Experience the difference in learning virtualization from the pioneers.

G.T. Enterprises has the credit of being the..... (In India) :-

- ☞ First VMware Partner since 1998.
- ☞ First VMware Authorised Consultants.
- ☞ First VMware Authorised Training Partner.
- ☞ First VMware Premier Partner.
- ☞ First to have received VMware Market Maker Award for APAC for two consecutive years 2005 & 2006.

The VMware Certified Training Program offers technology professionals the knowledge, skills, and credentials to deploy and maintain virtualization technology.

Benefits of VMware Professional Training :-

- ☞ VMware Infrastructure delivers measurable savings in both capital and operating costs.
- ☞ Demonstrate technical expertise to your customers.
- ☞ Enhance skills to sell, deploy and service VMware products.
- ☞ Ensure superior customer service and support of VMware products.
- ☞ Helps accelerate the application development and deployment lifecycles.
- ☞ Learn from VMware **Gurus** and enjoy the post training extended lab facility.

Seats are limited!! Hurry

We offer courses on VMware Virtualization 3 :-

- VI 3 : Overview - (1 Day).
- VI 3 : Install & Configure 3.5 - (4 Days) .
- VI 3 : Deploy, Secure & Analyze 3.5 - (4 Days).
- VI 3 : Virtual Desktop Manager - (4 Days).
- VI 3 : Fast Track 3.5 - (5 Days).

For more details contact us :-

Email : training@gte-india.com
Cell Phone : +91-9880223289, +91-9845227221
Tel : +91-80-26695890-94 (05 Lines), Fax : +91-80-26695887
URL : www.gte-india.com

G.T. Enterprises

"G.T.House", #48, 1st "B" Cross, 7th Block, Bhavani Layout, B.S.K. 3rd Stage,
Bangalore - 560 085, INDIA

TRAINING LOCATIONS :-

• Bangalore • Chennai • Gurgaon • Hyderabad • Mumbai • Mysore • New Delhi • Pune.



Make it Run. Anywhere, Anytime!

Did you ever dream that you wouldn't need to install applications? That you could just extract them into a folder and start using them? That you could even carry your set of applications, with all your data and settings, in a pen drive? Find out how...

It was Sunil's first day at work. And he spent the entire day with his systems administrator to set up the desktop he had been given, and get it running. He installed the complete software suite, including the mail client, Web browser, PDF reader, word processor, etc. It took him the rest of the day to configure *his* preferences for these applications, and download all of his mails in the new Mozilla Thunderbird. At the end of the day, a conversation over coffee went something like this:

Colleague: Hey Sunil, how was your day?

Sunil: Terrible! I just spent the whole day getting my system up and running. I had to configure everything again..."

Colleague: Yeah, I can understand... Had to do the same thing with my laptop—hard disk crashed a couple of days back. I mean, I can burn my songs and movies on a DVD, but what about applications?

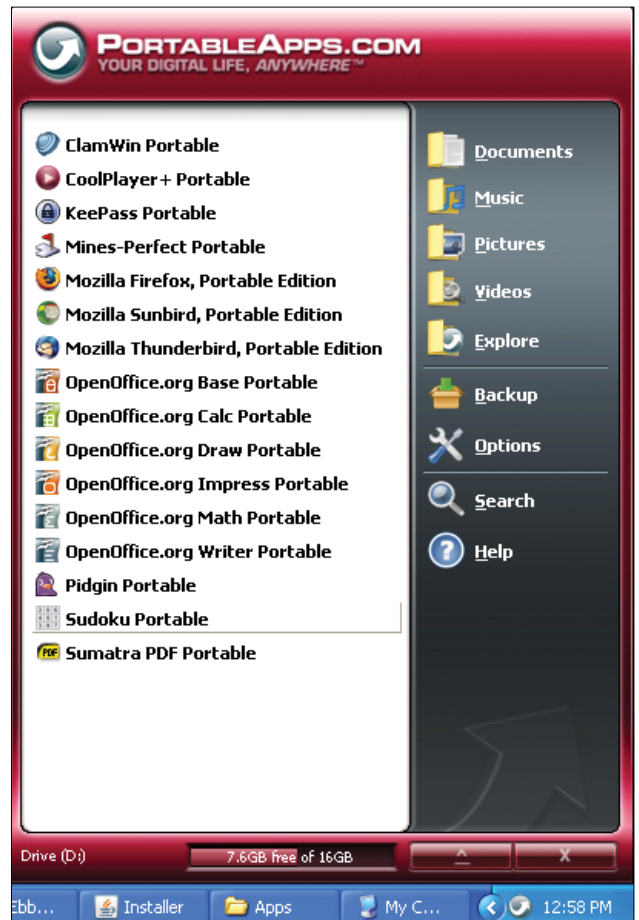
The solution to their problems lies in a typical 'Portable Application'.

True freedom...

What is a 'Portable Application'?

As the website www.portableapps.com says, "A portable app is a computer program that you can carry around with you on a portable device and use on any Windows computer." This is freedom from installing it multiple times.

Yes, download the application, install it (it is not the usual installation of software on Windows, which will create the registry entries; this 'installation' process is a simple extraction process), and there you have it. The installation happens 'in place'... or in a folder specified by you. All your settings, preferences, even saved passwords and browsing



Portable Apps Suite

history along with the application, will reside in *that* folder now. Just copy the folder onto a pen drive and carry it in your pocket...

Figures 1, 2 and 3 illustrate the steps taken during the installation of Mozilla Firefox.

Variety of applications...

Here is a short list of popular applications available in almost all areas

- Internet:
 1. Mozilla Firefox, Portable Edition with FireFTP extension
 2. Mozilla Thunderbird, Portable Edition
 3. Miranda IM Portable

WHO, WHEN AND WHY

PortableApps.com was founded by John T. Haller, the brain behind numerous portable applications, including Portable Firefox, which started off the portable software trend. The software has been made 'portable' since March 2004. It strives for a single, open platform usable by any software or hardware.

- Office utilities:
 1. OpenOffice.org Portable
 2. PDFTK Builder Portable (PDF viewer)
 3. GnuCash Portable
- Multimedia:
 1. GIMP Portable
 2. VLC Media Player Portable
 3. MPlayer Portable
- Games:
 1. PokerTH Portable
 2. Sudoku Portable
- Miscellaneous:
 1. 7-Zip Portable (file achiever utility)
 2. On-Screen Keyboard Portable
 3. SpeedCrunch Portable (calculator utility)

The complete list is available at portableapps.com/apps. Each of these applications has a forum and a support page associated with it. Also, the 'List of known issues' for each of the applications is well maintained. Likewise, you can report an issue or get queries answered.

Putting it all together: The PortableApps.com suite and platform

Everything put together, the 'PortableApps Suite' is a collection of commonly used applications bundled together in a single portable installation. You can download it from portableapps.com/suite. [The 'standard edition' of the suite is included in this month's *LFY* DVD].

Start contributing

The developer's forum here is pretty open minded and you can start contributing simply by registering yourself.

It is also a good idea for beginners to go through 'A (Not Finished) Beginner's Guide' at portableapps.com/development, which gives an overview of how to make an application portable.

Other than the 'stable' and 'working' builds in the Downloads section, there are several application builds available in alpha, beta and pre-release phases on portableapps.com/development/test. They include popular applications like a portable IRC client, XChat, the portable GVim, etc.

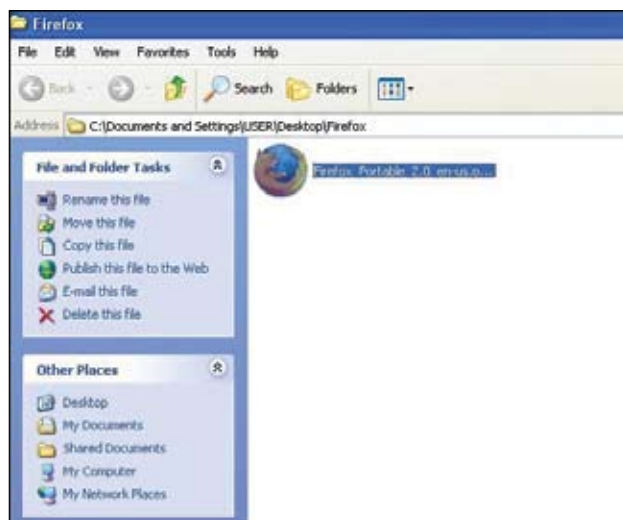


Figure 1: The portable Firefox app download that's ready for installation

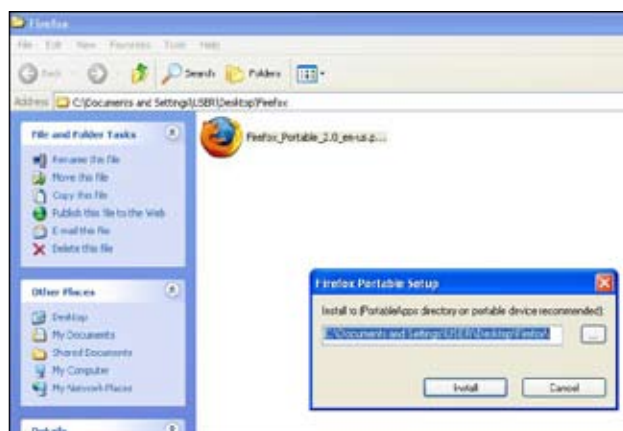


Figure 2: Specify the path where it should be extracted

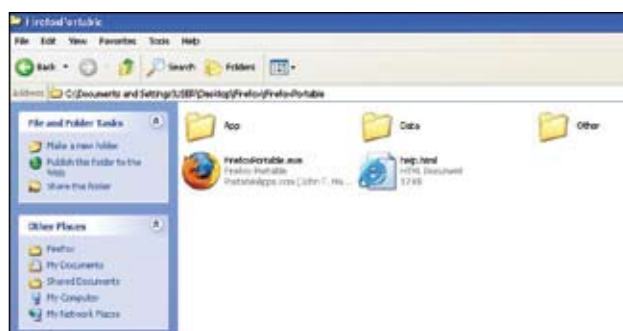



Figure 3: Ready to run!

Wait for it

The latest goodies, such as Google Chrome (that is, Chromium), are also in the process of being made portable. Even Linux applications are in line to be made portable. **END** 

By: Nilesh Govande. The author is a Linux enthusiast and could be contacted at nileshgovande@yahoo.com. His areas of interest include Linux system software, application development and virtualisation. He is currently working with the LSI Research & Development Centre, Pune.



Free(dom) Multimedia Powerhouse for Windows

From the rich family of Free Software in the Linux domain, here are the three most popular software for music and video, which are also available for the Windows platform.

I remember the old days when I was fully into Windows. At that time, installing software was quite a difficult task. Finding the correct version online and then installing it, unsure of whether it carried any adware, was a potential risk I used to take a lot, and paid the price for in terms of broken down systems. Another problem was the trial versions of software that would expire, like medicines, after a certain period of time. Once I *migrated* to Linux, these are issues I've never come across again.

There's so much free software available for Linux that you may get confused about which one is better. The added advantage is that most of these free software are available for the Windows platform as well. Here we are going to talk about three of the most *entertaining* free and open source

software for Windows. Using these open source software will also build confidence in you to make the choice of migrating to Linux—something I did a while back. (Now, I use Ubuntu!)

I am a big music and movies fanatic. I can't live without the two. So, a good movie player, an efficient music player and a really powerful media converter are some of the things that I need the most. We will start with the Songbird!

Songbird

Songbird is one powerful tool. You can find the software in the DVD accompanying this magazine. Installation is as easy as any other Windows software (Figure 1). Now, what can you *do* with this Songbird? Okay, let me rephrase the question here: what do you *want* to do? First, it is one of the most




Figure 1: Songbird installation




Figure 3: Libraries of your music in Songbird



Figure 2: Songbird default view




Figure 4: Songbird automatically recognises and connects your iPod

powerful music players (Figure 2). Once installed you can add media, the partitions or folders where your music files are saved, and Songbird will create a library of all of your songs (Figure 3) stored on your hard drive.

But these are all basic features that just about any music player gives us, and certainly not the reason why I like Songbird so much. I like it because it can double up as an efficient iPod manager tool as well as a full-fledged Web browser. We will first discuss the iPod manager. You need to install iPod support in Songbird. Simply click on *Tools*→*Add-ons*→*Install*. Now, browse the DVD and follow this path: *Software*→*multimedia*→*audio*→*Songbird*. Here, you will find two files; select the one that starts with the name *ipod-3*. When you click on this file, it will install iPod support and ask you to restart Songbird. Now, Songbird will recognise and connect your iPod automatically (Figure 4).

You can now rename any playlist or track; you can even edit the ID3 tags (Figure 5), which store meta data about your music, such as the album name, composer, song title, year and album art. Just like iTunes,



Figure 5: Editing ID3 tags in Songbird

you can transfer your songs and video to your iPod through the software. (I will later talk about how to easily convert your videos to be played on iPod.) Just create a playlist for the new album/song you want to transfer to your iPod, and drag and drop the album into that playlist.

Now, you have an option similar to iTunes, to sync music automatically or manually (Figure 6). When you sync the music, it will ask you whether your 'library has been synced with some other library' and whether you want to erase all the tracks on this iPod; click on 'No' (Figure 7), otherwise it will erase all the data, which you had earlier synced using iTunes. This is one stupid feature of iTunes that I hate the most. In the world of Linux,



Customised
Corporate
Trainings
Clustering,
Load
Balancing,
LDAP, PKI...

B
R
I
D
G
I
N
G

System
Integration
Staffing
Placements
Support

FOSTERing

Linux

RHCE
Shell
Scripting
PHP
Perl

T
H
E
G
A
P

RHCSS –
Clustering,
Mail Servers,
SysAdmin
Toolkit



redhat
TRAINING PARTNER

Exam Centre for RHCE & RHCSS

India's First RHCSS*

*Non-Red Hat Employee

FOSTERLinux
ing
Free & Open Source Training Education & Research

AWARDED MOST PROMISING CERTIFIED
TRAINING PARTNER 2007-2008

SCO- M22, II Floor
Old DLF Colony
Sector 14
Gurgaon
Tel: 0124-4268187
4080880,
9350640169,
9818478555

55, Deepak Bldg
13 Nehru Place
New Delhi - 19
Tel: 011-30880046,
30880047,
9810477448,
9891074905

Email: info@fosteringlinux.com
Website: www.fl.keenable.com



Figure 6: Syncing music with iPod




Figure 7: Warning while syncing music with iPod



Figure 8: Restore your iPod using Songbird

GTKPod and Amarok never do any such things. Not only do these two software enable you to simply drag, drop and save your files on your iPod, but they also allow you to take back-ups of music from the iPod onto your PC. But that is possible only if you are using Linux. So, I wish you migrate to GNU/Linux soon.

Songbird has another important feature: restoring your iPod if it ever gets corrupted. Click on the iPod icon and select the *Tools* tab next to the *Music* tab to restore to factory settings.

Coming back to the music features, Songbird has a full arsenal of add-ons or extensions. Connect to the Internet and check the option 'add-ons' on the default page of Songbird (Figure 9). Click on this option and it will open a page of



Figure 9: Add-ons option in Songbird



Figure 10: Lyrics manager in Songbird




Figure 11: Ready the lyrics of the currently playing song

add-ons for you.

One of the most interesting features you will find is *Lyrics* (Figure 10). Click on this and it will install a script, which will enable you to automatically connect to the Internet and show you the lyrics of the song that you are currently playing.

You can also save those lyrics locally so that the next time you are listening to that song, you don't have to connect to the Internet to view the same lyrics (Figure 11).

You can also listen to Last.fm and share music with your friends. As I said earlier, Songbird doubles up as a Web browser also. Click *Ctrl+T* and you will see a new tab, just like in Firefox, and you can start browsing the Internet; check your e-mails if you like.



Figure 12: Choose format and save location




Figure 13: Pre-sets available in VLC



Figure 14: VLC converting video

VLC

VLC is one of the most popular and powerful movie players and converters. You will find it in the accompanying DVD. Installation is simple. Once you install VLC, you will not need to install any other codec or software to play different movie formats from MPEG to DIVX. Now, we are going to talk a bit about the converter feature of this player.

Click on the *Media* menu in VLC and look for *Convert/Save* in the drop-down box. Click on it and it will ask you to choose the file you want to convert. Once you select the file, you will be greeted by the next window (Figure 4: Choose format and save location). This one is important. Here you will have to choose the format you want to convert your video into, and the



Figure 15: Avidemux default window



Figure 17: Selecting the audio format




Figure 16: Selected file on preview window

place you want to save it.

Now, you can also use some pre-sets for converting video for various devices including the iPod, Xbox and many more (Figure 13). Now, click on *File* and use *Browse* to select the location you want to save the converted file. Then click *Save* and your file will start converting (Figure 14). It may take a while, depending on the size of the file and the hardware specifications of your PC.


Avidemux

Avidemux is one of the most popular media converter software on Linux machines. It is available for Windows users as well. You can get the software in the current *LFY DVD*. The installation is, as usual, pretty simple. (Just keep clicking those *Next* buttons! I still find software installation on Linux to be much easier—not so many ‘click-nexts’! You can try that once you migrate to Linux. You already have the Mandriva 2009 DVD accompanying this magazine, why don't you give it a try? And don't blame me if you also start using Linux from there on. It's addictive; trust me!).

Now, you can find an Avidemux

icon on your desktop; Figure 15 shows the default window after launching the program. Here, you can see several options. Select the file you want to convert and you will see it in the preview window (Figure 16).

The next step is to choose the format you want to convert your video into. Here, we will tell you how to convert your video into a format playable on your iPod. For selecting the video format, click on the *Copy* option under the *Video* section (Figure 16) and choose the format you want to convert your video to. In case of the iPod, select MPEG-4 ASP (laavc). Now, select the audio format (Figure 17); for an iPod, we will select AAC (FAAC). The next step is to select the format of the file, so check MP4. Now, click on *Save* and select the location you want the video to be saved. While giving the name of the file that it should be saved as, you will have to add the corresponding extension, in this case *xxx.mp4*. Now, click on *Save* and you will have your video saved in the desired format.

Avidemux also allows you to edit your videos. You can crop the screen size, increase/decrease brightness and contrast, and even add subtitles. There are many such filters that make Avidemux a very powerful converter. Well, Avidemux is installed on your PC, right? So please start playing with it because I've run out of the allotted word-limit! **END** 

By: Swapnil Bhartiya, assistant editor, EFYTimes.com

ATTENTION
MCA, MSC (IT)
Final Year Students
GET Trained
Project Placed 

www.lynusacademy.com/projects

- Projects available in LAMP, Java, J2EE, Perl, Python & ROR
- Batch Discount available
- Accommodation will assist For Out Station Students
- Register Now .. Limited Seats Available
- IEEE Projects also available

 Authorized Training Partner

 **redhat**
CERTIFIED TRAINING PARTNER

Special Offer for
Job hunters & Freshers
Zend PHP Training
(3 Months)
Rs. 23,200 /- Only

Other courses

- RHCT & RHCE
- Shell Scripting
- Perl Scripting (for Sys. Admin)
- G-FOSS - Web Development (3 Months & 6 Months)
- G-FOSS - Engineer (3 Months & 6 Months)
- Perl + CGI Programming
- Python Programming
- Ruby on Rails
- AJAX & JQuery
- MySQL - for Developer
- MySQL - For Administrator
- PostgreSQL - for Developer
- PostgreSQL - for Administrator



Future is open

Lynus Academy Pvt Ltd


Chennai - 044-25021630,
42171278,9840880558

Pondicherry - 0413-2276233,
4202121,9952966528

contactus@lynusacademy.com
www.lynusacademy.com



Non-Greasy Ways to **CHANGE** the Flat Tyres of Your Car



Have you ever got stuck with a flat tyre? We all have. So now, you're going to hear about how to change flat tyres of your car. If you are a businessman for whom every rupee matters, this article is going to save you a few lakhs, if not more. And even if you are just another common man or woman, you're still going to save some money.

I just changed my mind and instead of telling you how to change all the four tyres of your car, I am going to tell you about something else. But I first need to know if you use computers. Of course you do—that's why you are reading this article. And there may be chances that you use Microsoft Windows. This issue is all about how to use free and open source programs on

Windows. So, it is likely (even remotely) that you are running Windows on your PC. But you don't really use Windows. You use applications to watch videos, listen to music, view and edit photographs, and to browse the Internet and check your e-mails. I guess these are the four things that you use your computer for. Ah! there can be six or even more things, but we won't get into that. What we will get into is how to save some of your

money and also ensure that your PC is free of adware and Trojans. Most importantly, we bring you a fresh air of 'free' software.

One step at a time!

When was the last time you logged into your e-mail account or browsed the Internet? What? Three years ago? Do it a bit more frequently! I think you may be using Internet Explorer or something with a cute *e* for a logo. It can do a few things, but it can't do a lot of things. Then what performs 'all functions'? You should try the DVD that you got with this magazine; look for Firefox and install it. Once you've done that, I will tell you all that Firefox can do. So, come back once it's on your machine.

I know you didn't go anywhere, so let's continue. Firefox is like the 'Windows' of the online world. It not only offers you the basic functionality of browsing the Web and allowing you to check your e-mails (do you have an e-mail account?), but also opens a new world of online tools. Just like you have applications in Windows, there are plug-ins or add-ons in Firefox. There are thousands (I am not sure about the exact numbers; was never good at maths) of such extensions that can do a lot for you. I would have covered all the extensions, but you know I have to go and change the flat tyres of my car. So, I will leave you with some glimpses of Firefox. There are extensions that will download videos from services like YouTube and Google Video, and can convert online videos in the format you want. So install Firefox and then check it out!

Read between the lines

If you are reading this article, I guess you use MS Word to type letters, create presentations with PowerPoint, use Excel and many such other tools. How much did it cost you? What if I tell you that there is free petrol... no, I mean to say, what if I tell you there is a program that is not only very powerful and effective but is also available for free? This software is not only used by millions of individuals but also by many organisations and governments. It is called OpenOffice.org, or OOo in short. Once again, go to your PC and find the software in the DVD that came with this magazine and install it. OOo installs in a few minutes and can do almost everything you want to with a word processor (Word), spreadsheet (Excel) and presentation software (PowerPoint). Well, I'm currently writing this article on OpenOffice.org Writer. So, next time when you write feedback to the magazine, use OOo Writer.

Why so serious?


Are we getting too serious? Should I tell you a joke? *No?* Then let's talk about the movies. Do you watch movies or listen to songs? Oh! That was a bad question. Yes, everyone does watch movies. Where do

you watch them? I don't mean at which theatre, but where *on your PC*. I think either it's Windows Media Player, or iTunes and QuickTime. They are good, but nothing is perfect. Try VLC and Songbird. Two things about each of these. VLC not only plays almost any format of movies, but also convert movies from one format to another. Songbird is not just a music player like Media Player or iTunes, but can also double as a browser, and sync music and video with your iPod. We will discuss that in a separate article. So, now you have new mates—VLC and Songbird—for Media Player that comes with Microsoft Windows.

Trigger happy!

I had a friend who was great at shooting. But he got killed during one shoot. He had run out of bullets. So, do you shoot? And if so, do you put your photos online at Flickr, Piccasa, or some social networking site? Do you edit your pictures before you put them online—increase brightness, adjust contrast, etc? Or you might even have cut your own image and pasted it in the foreground of the Taj Mahal to impress friends—I did that once. How did you do that? Did you use Photoshop? Don't tell me that you used Paint Brush—it can't do much. But Photoshop is too expensive a tool to be used for cutting your image and putting it against one of the wonders of the world. Don't do that again. If your life depends on photo-editing tools, and you use them a lot for basic as well as advanced image manipulation, I suggest you neither waste money nor compromise with such a very basic tool. I will once again ask you to go and check the *LFY* DVD. There you will find something called The GIMP. Install it! It is one of the most powerful free software image editors I have ever come across. If you know Photoshop, you will find an alternate for almost every feature of Photoshop in The GIMP. Now, go ahead and try cutting your image and put it against the Pyramids of Egypt.

Well, there is a lot more I can tell you about other replacements or ways to put FOSS tools onto Windows. But since this article is *not* about putting FOSS tools on Windows, but rather about '*Non-Greasy Ways to Change the Flat Tyres of Your Car*', I am not going to tell you that. You can skip to the next article to learn all those FOSS tools that run on Windows, which you are getting with the DVD accompanying this issue.

But did I hear you say this article is *not* about '*Non-Greasy Ways to Change the Flat Tyres of Your Car*'? Well, you may have missed the point, then. Don't you think we replaced all the four tyres of the car called Windows? Hope to see you driving a Linux car next time. ;-) 

By: Swapnil Bhartiya, assistant editor, *EFYTimes.com*



Free Software for Windows

THE TOP
25

One of the DVDs accompanying this magazine contains more than 200 Free Software application for the Windows platform. In the following few pages we have highlighted some of the bests amongst them. Of course, this list excludes the nifty multimedia tools that we have covered in an article titled “Free(dom) Multimedia Powerhouse for Windows” on Page 34. Heck, still wasting your time reading this intro? Go check out the various software. We hope you have enough space in your hard drive. :-)



Internet

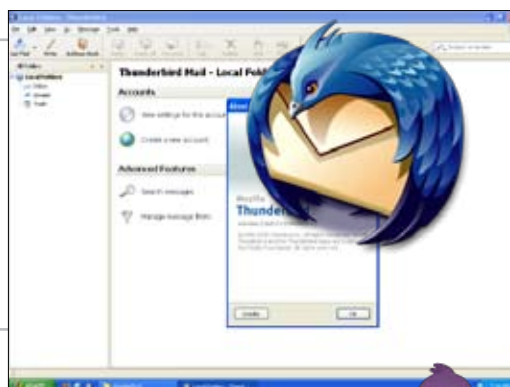
Mozilla Firefox

Mozilla Firefox is considered to be the second-most widely used browser worldwide, after Internet Explorer. Firefox uses the Gecko layout engine to display Web pages. It includes tabbed browsing, a spell checker, incremental find, live bookmarking, a download manager, and an integrated search system that uses the search engine you want it to. Firefox 3 comes with dozens of new features, including the smart location bar, one-click bookmarking and blindingly fast performance. It integrates elegantly with the installed anti-virus software so that when you download a file, your computer's anti-virus program automatically checks it to protect you against viruses and other malware, which could otherwise damage your computer. This feature is available only in the Windows version of the browser.



Mozilla Thunderbird

Mozilla Thunderbird makes e-mailing safer, faster, and easier with such features as intelligent spam filters, a built-in RSS reader, and quick search. Thunderbird was designed to prevent viruses and to stop junk mail so you can get back to reading your mail. You can automatically have your junk mail deleted or you can put it in a folder you specify, just in case you like reading junk mail.



Pidgin

Pidgin is a multi-protocol instant messaging client. It includes support for AIM, ICQ, Jabber/XMPP, MSN Messenger, Yahoo!, Bonjour, Gadu-Gadu, IRC, etc. Pidgin can log in to multiple accounts on multiple IM networks, simultaneously. Pidgin features some of the standard tools for an instant messaging client, such as tabbed conversations, a contact list, file transfer on supported protocols, as well as conversation and chat logging. Tabbed conversations are an optional feature on Pidgin. The IM window consists of the message window, formatting tools, and an edit box.





Internet

PuTTY


PuTTY is a terminal emulator application that can act as a client for the SSH, Telnet, rlogin, and raw TCP computing protocols. The name 'PuTTY' has no definitive meaning, though 'tty' is the name for a terminal in the UNIX tradition, usually held to be short for 'teletype'.



WinSCP


WinSCP (Windows Secure copy) is an open source free SFTP client and FTP client for Windows.

Legacy SCP protocol is also supported. Its main function is the safe copying of files between a local Windows and a remote Linux computer. WinSCP can act transparently as a remote editor. When the user clicks on a (text) file in the remote file manager, it transfers the file to the local machine and opens it in the integrated editor, where Windows users can feel very much at home. WinSCP also offers basic file manager functionality. For secure transfers, it uses Secure Shell (SSH) and supports the SCP protocol in addition to SFTP.



FileZilla

In today's fast paced world of the Internet, we often need to connect to a FTP site and upload/download files from there. FileZilla Client is a fast and reliable cross-platform FTP, FTPS and SFTP client with lots of useful features and an intuitive graphical user interface. It supports the resume and transfer of large files, even greater than 4 GB.



HTTrack

Have you ever thought of browsing a website without an Internet connection? Or storing all the information of the website for offline use? HTTrack, a free and easy-to-use offline browser utility, addresses both those needs. It allows you to download a World Wide Web site from the Internet to a local directory, recursively building all directories, getting



Inside DVD

Browsers

- Amaya Web Browser
- Arora
- Firefox 3.0.3
- Flock
- K-Meleon 1.5.0
- SeaMonkey 1.1.12

IM Tools

- AMSN
- Ayttn 0.4.6
- Coccinella 0.96.10
- CSpace 1.27
- Instantbird 0.1.2.1
- Miranda IM 0.7.10
- Pidgin 2.5.1
- Psi
- Sim-IM
- Spark 2.5.8
- WengoPhone 2.1.2

Mail Clients

- Evolution
- Mozilla Thunderbird 2

Internet Utilities

- Ares
- Azureus
- BitTorrent 6.1.1
- DeHinter 0.1
- FileZilla Client 3.1.3.1
- FreePOPs 0.2.7
- FrostWire 4.17.0
- Galet 0.3
- Getleft
- HTTrack
- LimeWire 4.18.8
- Magic Mail Monitor (MMM)
- POPFile 1.0.1
- PopMan 1.3.0
- Putty
- RSS Bandit
- RSSOwl
- Shareaza 2.4
- Soulseek
- WinSCP 4.1.7
- wxDownload Fast



Business Applications


SugarCRM

SugarCRM is one of the most popular CRM software. It can cater to the requirements of companies of any sizes. Sugar enables organisations to efficiently organise, populate and maintain information on all aspects of their customer relationships. The system also offers a graphical dashboard to track the sales pipeline, the most successful lead sources, and the month-by-month outcomes for opportunities in the pipeline. This CRM software can be customised to the needs of any business and can be installed on systems running Windows. It offers a more flexible, cost-effective alternative to proprietary applications. SugarCRM offers several deployment options, including on-demand, on-premise and appliance-based solutions to suit customers' security, integration and configuration needs.



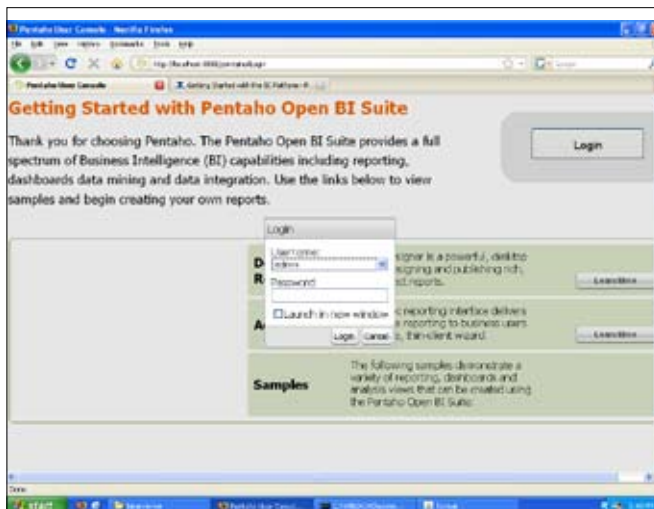
Adempiere

Adempiere is an open source software solution that includes enterprise resource planning (ERP), customer relationship management (CRM) and supply chain management functionality for business processes. It is a community-based project that can be installed on Linux as well as on Windows. Adempiere provides a framework for extending and customisation so that it can meet your business needs.



Pentaho BI Suite

Every business has lots of data and information stored. The collection, integration, analysis, and presentation of business information are termed as business intelligence (BI). The Pentaho BI Suite provides a full spectrum of business intelligence (BI) capabilities including query and reporting, interactive analysis, dashboards, ETL/data integration, data mining and a BI platform. Pentaho is one of the world's most popular open source BI suites that can be installed on Windows/Linux. The platform includes an embedded workflow engine and can be easily integrated into business processes. The framework provides core services including authentication, logging, auditing, workflow, Web services and rules engines.



Inside DVD

ERP and CRM

- Abstract Accounting ERP
- Adempiere
- Commander4j
- Compiere ERP + CRM Business Solution
- Dolibarr
- JFire
- Neogia
- OpenBlueLab
- Openbravo ERP
- openCRX
- Plasma
- SplendidCRM Open-Source
- SugarCRM
- Tiny ERP
- vtiger

Business Intelligence:

- Openproj
- Pentaho BI Platform
- SpagoBI

Accounting

- Buddi
- GFP
- GnuCash
- GRISBI for Windos
- Money Manager Ex
- OsFinancials
- TurboCASH



Office Productivity

Inside DVD

Office Tools

- AbiWord 2.6.4
- Gnumeric
- OpenOffice.org 3.0.0
- PDFCreator
- PDFsam (PDF Split and Merge)
- Scribus 1.3.3.12
- SumatraPDF

Compressing/ZIP Tools

- 7-Zip
- Jar Ajar
- PeaZip -- The free Archiver

Utilities

- BwgBurn
- Clamwin
- InfraRecorder
- KeePass
- PortableApps
- Password Safe
- regain
- VirtuaWin
- WinMerge

PDFCreator

PDFCreator is an open source tool to create PDF files from nearly any Windows application. It can be used to create PDFs from any program that is able to print. PDFCreator allows you to disable printing, copying of text or images and modifying the original document. The user can also choose between two types of passwords, user and owner, to restrict access to PDF files in several ways. Besides being installed as a printer driver, PDFCreator can be associated with .ps files to manually convert PostScript to PDF format.



Gnumeric spreadsheet

Gnumeric spreadsheet is part of the GNOME, a project to create a free, user friendly desktop environment. The goal of Gnumeric is to be the best possible spreadsheet. It can read files saved with other spreadsheets. Gnumeric handles large spreadsheets while remaining responsive.



OpenOffice.org

OpenOffice.org is a free cross-platform office application suite for word processing, spreadsheets, presentations, graphics, databases and more. It is available for a number of different computer operating systems. It supports the ISO-approved Open Document Format (ODF) for data interchange as its default file format, as well as Microsoft Office 97-2003 formats, among others. OpenOffice.org is the best alternative office suite available from the free software world. After three years of continuous improvement, OpenOffice.org has currently reached the landmark version 3.0. This includes support for the Microsoft Office 2007 format, improved crop feature in Draw and Impress, along with spreadsheet collaboration through workbook sharing, etc.





Multimedia Apps


Audacity

Audacity is a free, easy-to-use audio editor and recorder for Windows, Mac OS X, GNU/Linux and other operating systems. Audacity can be used to record live audio; convert tapes and records into digital recordings or CDs; and to edit Ogg Vorbis, MP3, WAV or AIFF sound files. Audacity can record live audio through a microphone or mixer, or digitise recordings from cassette tapes, vinyl records, or mini discs. With some sound cards, it can also capture streaming audio.



Mp3SpliT

Mp3SpliT is a utility to split MP3 and Ogg files by selecting a start and an end time position, without decoding. It's very useful to split large MP3/Ogg to make smaller files or to split entire albums to obtain original tracks. If you want to split an album, you can select split points and filenames manually, or you can get them automatically from CDDb (the Internet or a local file) or from .cue files. It also supports automatic silence split that can be used to adjust CDDb/cue split-points. You can extract tracks from MP3Wrap or AlbumWrap files in a few seconds.



DVDStyler

There are many DVD authoring tools available for Windows. DVDStyler is a cross-platform DVD authoring application that makes it possible for video enthusiasts to create professional-looking DVDs. This stands out from the rest of the tools as it is free and open source. It supports AVI, MPEG and VOB files, and allows putting files with different audio/video formats on one DVD. DVDStyler can create a DVD video with interactive menus and can import image files for the background. It allows you to place buttons, text, images and graphic objects anywhere on the menu screen.



Inside DVD

Audio Players

- Aqualung
- aTunes
- Cactus Jukebox
- CoolPlayer
- Jajuk
- Mixxx
- musikCube 1.0
- SnackAmp
- Songbird
- Zinf Audio Player

Video Players

- Kantaris Media Player
- Miro 1.2.7
- Media Player Classic
- MPlayer
- VLC-0.9.2

Audio Editing Tools

- Audacity
- AudiobookCutter
- BonkEnc
- Cdex
- MediaCoder
- Mp3split
- SoX
- Streamripper

Video Processing Tools

- Avidemux
- DVDStyler 1.7.0
- Krut Computer Recorder
- MediaInfo
- Movica
- OrDrumbox
- StreamHijacker
- VirtualDub
- YAAI



Graphics

Blender

Blender is a free and open source 3D graphics software. It can be used for modeling, UV unwrapping, texturing, rigging, water simulations, skinning, animating, rendering, particle and other simulations, non-linear editing, compositing, and creating interactive 3D applications. Blender is available for several operating systems. It has a feature set somewhat similar to other high-end 3D software such as Cinema 4D, Lightwave and Maya. Blender features an internal filesystem that allows you to pack multiple scenes into a single file called a .blend file. All of blender's .blend files are forward, backward, and cross-platform compatible with other versions of Blender, and can be used as a library to borrow pre-made content.



Inside DVD

- AlbumShaper 2.1
- Blender 2.48
- Dia
- Gallery Mage
- GIMP 2.6
- GLIPS Graffiti Editor
- GPixPod
- Inkscape
- K-3D
- Paint.NET 3.36
- Sweet Home 3D 1.4

FunStuff Inside DVD

FUNSTUFF

- Bifusion
- Celestia
- Crack Attack!
- Enigma
- FloboPuyo 0.20
- FreeCol
- Frets on Fire
- IceBreaker
- Kobo Deluxe
- LBreakout2
- Pang Zero
- Scorched3D
- Sokoban YASC
- SuperTux-0.1.3
- TORCS
- Tuxtyping
- World Wind 1.4
- X-Moto

Inkscape

Inkscape is an open source graphic editing tool that uses the W3C's scalable vector graphics format (SVG), and can be used as an alternative to Illustrator and CorelDraw. Some supported SVG features include basic shapes, paths, text, markers, clones, alpha blending, transforms, gradients, and grouping. In addition, Inkscape supports Creative Commons' metadata, node-editing, layers, complex path operations, text-on-path, text-in-shape, and SVG XML editing. It also imports several formats like EPS, PostScript, JPEG, PNG, BMP, and TIFF, and it exports PNG as well as multiple vector-based formats.



The GIMP

The GIMP (GNU Image Manipulation Program) is a freely distributed program for tasks such as photo retouching, image composition and image authoring. It can be used as a simple paint program, an expert quality photo retouching program, an online batch processing system, a mass production image renderer, or as an image format converter. The GIMP can also be used to create basic animated images in GIF format. It is often used as a free alternative software replacement for Adobe Photoshop.





Developers


WAMP server

The idea of developing and hosting a website leads us to think about the cost involved in purchasing different components, apart from the cost of Windows. WAMP (Window, Apache, MySQL, PHP/Perl/Python) is a free alternative. The components of WAMP are available as free/open source software. This means that a dynamically-generated website can be set up without major software purchases or update subscription costs. WAMP is a kind of mini-server that can run on almost any Windows version. WAMP includes Apache 2, PHP 5, and MySQL (phpMyAdmin and SQLitemanager are installed to manage your databases) pre-installed. It can be easily installed on the local system to test and develop a full-featured website.



NetBeans

An IDE is a software application that provides comprehensive facilities to computer programmers for software development. IDEs are designed to maximise programmer productivity. NetBeans is a free, open-source IDE for software developers. It consists of all the tools used to create professional desktop, enterprise, Web, and mobile applications with the Java language, C/C++, and Ruby. The NetBeans IDE is easy to install and use straight out of the box and runs on many platforms including Windows, Linux, Mac OS X and Solaris.



Python

Python is a clear and powerful object-oriented programming language, comparable to Perl, Ruby, or Java. Python supports multiple programming paradigms (primarily object oriented, imperative, and functional) and features a fully dynamic type of system and automatic memory management. Python is often used as a scripting language and can be easily extended by adding new modules implemented in a compiled language such as C or C++. It can also be embedded into an application to provide a programmable interface. This general-purpose, high-level programming language runs on many platforms, including Windows.



Inside DVD

Databases

- Berkley DB 4.7.25
- FireBird
- MySQL 5.0
- PostgreSQL
- SQLite

IDE's

- CodeBlocks
- CodeLite
- Dev C++
- DrPython
- Eclipse
- Jedit
- JUDO
- Lazarus
- Netbeans 6.1
- SharpDevelop

Compilers and Assemblers

- FreeBASIC Compiler
- NASM
- Free Pascal Compiler
- Python 2.6
- Ruby
- WinTclTk

Electronic Design & Automation

- Kicad EDA
- LayoutEditor
- Logisim
- Quite Universal Circuit Simulator (QUCS)
- WinAVR
- XCCircuit 3.4.10

Web Development Tools

- AppServ
- NVU
- PHPMyadmin
- WAMP Server
- XAMPP

Installer Package Creators

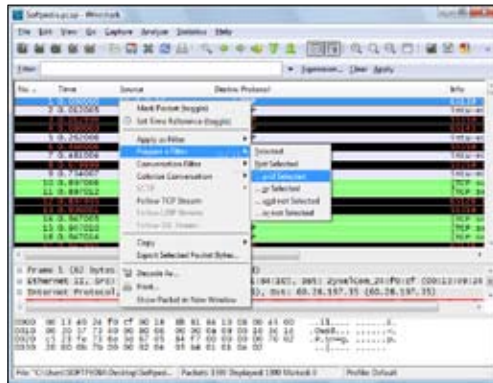
- Appupdater
- Inno Setup Unpacker
- NSIS (Nullsoft Scriptable Install System)



Power Users

Wireshark

Wireshark is the world's foremost network protocol analyser. This packet sniffer is used for network troubleshooting and analysis. Wireshark has a rich feature set that includes deep inspection of hundreds of protocols, with more being added all the time. It also has live capture and offline analysis and a standard three-pane packet browser. It uses the cross-platform GTK+ widget toolkit and thus can run on Windows as well.



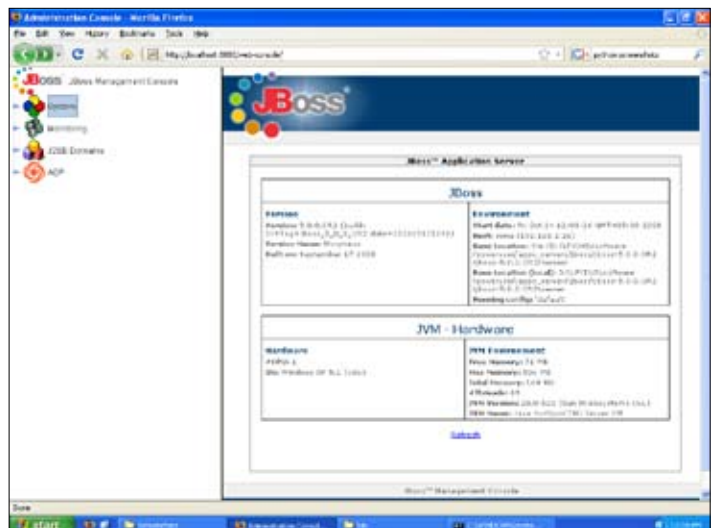
QEMU

QEMU is a generic and open source machine emulator and virtualiser. When used as a machine emulator, QEMU can run OSs and programs made for one machine (for example, an ARM board) on a different machine (for example, your own PC). By using dynamic translation, it achieves a very good performance. When used as a virtualiser, QEMU achieves near-native performances by executing the guest code directly on the host CPU.



Jboss Application Server

Jboss Application Server (or JBoss AS) is an open source Java EE-based application server. Because it is Java-based, the JBoss application server is cross-platform and can be used on any operating system that Java supports. It is the most popularly used application server on the market.



Inside DVD

Application Server

- Apache Geronimo v2.1.3
- JbOSS
- JOnAS
- Zope

Network Management tools

- AutoScan-Network
- Nmap
- Packetizer
- TightVNC
- UltraVNC
- Wireshark

Backup and System Management Tools

- Areca Backup
- Bacula
- Cobian Backup
- Daphne 1.36
- Eraser
- Explore2fs
- Magical Jelly Bean Keyfinder
- Startup Manager
- TestDisk 6.10

Virtualization

- Bochs
- coLinux
- DSL
- QEMU



Run Linux



'on' Windows in

What if I tell you that you can install Linux 'in' Windows just like any other application? This article will guide you on how to use Ubuntu Linux without messing up with your partitions and without worrying that you don't know too much about Linux. It is as easy as sending out an e-mail.

Are you a Windows user who has always been apprehensive of using Linux, as many people say it's too complicated for home users? I guess they are correct, because it's quite clear that they have not touched any Linux machine for at least the last few years.

Linux is today much more simple to install on any machine, but this is not what we are going to discuss. We are going to tell you about something even more exciting. But before we do that, let me ask you if you have ever installed any software on your Windows machine. And if you have, how difficult was that? Like taking a walk in the park, right?

Now, what if we tell you that you can now install and run one of the most popular Linux operating systems, Ubuntu, inside your Windows machine just the way you install any other software.

Let me repeat: You can now install Ubuntu in your Windows just as any other

software and you can very easily remove it too, again, like another software. So, if you are interested, let's take the next step.

The Hardy Heron (Ubuntu 7.04 LTS)

What you need to start with is a LiveCD of the Ubuntu operating system. The Intrepid Ibex, a.k.a Ubuntu 7.10, will be released on October 30, and the *LFY* December issue will bundle that. Turn on your Windows machine and insert the Ubuntu CD in the drive. In a few seconds you will see a window pop up (Figure 1). You will find three options here. Select the second one that says, 'Install Inside Windows'. When you click on this, you will be greeted by another window (Figure 2) that says you are about to install Ubuntu. You can see six options here.

Option 1. The 'Installation drive' is the partition on which you want to install Ubuntu. You can change the partition by selecting it from the drop-down menu. Choose a partition that has at least 5-10 GB of free space.

The second option says 'Installation size'. This is the space on your hard drive that you want Ubuntu to use. We recommend you give at least 6-10 GB of space to Ubuntu. Now you can select that too from the drop-down menu.

Leave the third option as it is. The fourth option is selecting the language, which you should leave as it is, unless some other language interests you in the drop-down menu.

The fifth option is the username. By default, your username will be taken from your Windows account, but you can change that if you wish, and the last option is about the password. Select these two and remember it, as you will be needing these to log in to Ubuntu once it is installed.

Once all options have been selected, you can click on *Install* and your installation will start. At this stage it will basically copy and install some important files on your machine. Once the copying is being done (Figure 3 and 4), you will be greeted by the screen shown in Figure 5, where you are asked to restart your machine. This will also open the DVD/CD drive automatically and you must remove the LiveCD from the drive.

Now, when you restart your system, you will see a

black screen with some text on it. Don't panic or worry. Let it do whatever it has to. Then you will see the text shown in Figure 6. Now, use your keyboard's up-down arrow keys to select Ubuntu and hit *Enter*. This will take




Figure 3: Windows checking Linux installation files



Figure 1: Install Linux just like another window application




Figure 4: Copying image on hard drive

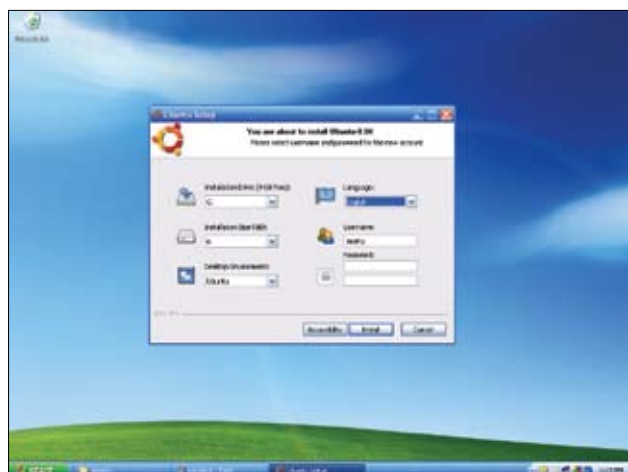


Figure 2: Selecting drive to install Linux on




Figure 5: Restart your pc

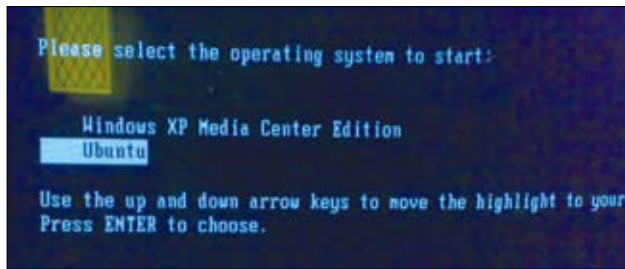


Figure 6: Select to Ubuntu




Figure 7: Final installation of Ubuntu



Figure 8: Ubuntu login screen

a few moments, following which you will see Ubuntu on your screen. Wait for a while, the next screen will be an Ubuntu wallpaper and you will see a box (Figure 7) where Ubuntu will install the operating system and prepare your machine. This can take around 10-15 minutes depending on your system specifications.

Once the installation is done, the PC will again restart, and once again you will be greeted by the same screen shown in Figure 6. Again, select Ubuntu and now you are in! This time, you will be greeted by a typical Ubuntu log-in window (Figure 8), where you will have to enter your username and password—hope you still remember them.

Now, you have Ubuntu Linux running on the same machine (Figure 9). So, was it difficult?

Something important now: you will be able to see other partitions under *Places--> Computer*, located on the top menu bar of your Ubuntu screen. Remember, if you want to save any data, save it *only* on these partitions that are also accessible through Windows. If you save anything on the Ubuntu desktop, you will not be




Figure 9: Finally you are running Ubuntu




Figure 10: Use 'Add/Remove Programs' to uninstall Ubuntu

able to see it in Windows.

I understand the user interface of Ubuntu is quite different from what you're used to in Windows, but hey, it's just a matter of time before you get used to this one too.

You can read more about how to install software in Ubuntu and many other things at help.ubuntu.com. The URL has everything that a new user requires, written in plain text. So, what are you waiting for; get started with Linux.

Uninstall

If you plan to uninstall Ubuntu from Windows because you want to go for an independent Ubuntu installation, go to the Windows Control Panel and under the *Add/Remove software* section, you can find Ubuntu. Uninstall it the way you uninstall any other software (Figure 10). But remember, save all your data only on the other partitions of Windows, so that you can use that data even after you have completely removed Ubuntu from your machine. **END**

By: Swapnil Bhartiya, assistant editor, *EFYTimes.com*,
and **Niraj Sahay**, *LFY bureau*

Microsoft®

EQUIP YOURSELF

Microsoft invites all developers to join the open source community and online resource centre. Interact with the likes of Sam Ramji, Director of Open Source Software Lab, Jamie Cannon, Lead of OSS Community & Platforms and many others to find solutions to all interoperability challenges.

JOIN WWW.PORT25.IN, A MICROSOFT PLATFORM FOR OPEN SOURCE DEVELOPERS.

PORT25



McCam/MSFT/08



Schedule Your Tasks

It's very easy with *cron*—so let's get started.

Have you ever wished to run applications on given schedules? Consider the following issue. Some ISPs nowadays provide free Internet usage during off-peak hours, i.e., after midnight. For example, BSNL DataOne broadband users enjoy free downloads between 2 a.m. and 8 a.m. But staying sleepless to accomplish those free downloads seems quite a strain. What if you missed getting up at 8 a.m. and the downloads ate up your bandwidth and money? Scheduling your downloads within free bandwidth hours could be a way out of these issues. But, how do we achieve this?

Linux administrators (or even those geek users) very often need to execute some programs on a regular basis. For example, the admin might need to monitor the disk usage of a system. His best bet in most cases would be *cron*, a handy solution to execute several tasks at a given time schedule. It is a utility written by Brian Kernighan, made available from UNIX version 7. Let us dig into this classic UNIX tool and find out how it can schedule our downloads too.

Consider some of the scenarios when you can use *cron*:


- If you run a website that deals with a large number of images and you want to create thumbnails automatically during a given time period every day or week.
- If you want to keep track of your back-ups synchronised easily without much pain and effort.
- Most importantly, if you want to run file downloads and torrents in a specified time.
- If scheduling of automatic system updates is required.

To define *cron* in a classical way, it is a daemon that runs in the background as a service. In order to create scheduled jobs, we use the command *crontab*.

How to create a scheduled job

Open a terminal and enter the following command:

```
crontab -l
```



It will list the currently installed crontable:
To edit the list of jobs in *cron*, you can run:


```
crontab -e
```

It will open the default text editor to let you manipulate the jobs. After you are done making changes, save and exit the editor. It will immediately activate all your *cron* jobs.

Or alternatively, open the terminal and enter the following command:

```
gedit joblist.cron
```

This will launch the *gedit* text editor (replace *gedit*



with *kwrite* if you use KDE) after creating a file called *joblist.cron*. Enter the following line in the file:

```
00 * * * * date >>/tmp/hourlytime.log
```

Save and exit. Now, back in the terminal, run the following commands:

```
crontab joblist.cron
```

...followed by:

```
crontab -l
```

This will now list the following:



Figure 1: crontab syntax

```
00 * * * * date >> /tmp/hourlytime.log
```

I'm sure you must be wondering how to read the above line. There is a special syntax in which *cron* jobs are to be entered. Figure 1 explains it.

Environment settings

You might have thought that since *cron* runs on a shell, it will use the same environment settings as the parent shell. But that is not true. We have to specify separate environment variables by adding them to *crontab*. For example:

```
DISPLAY=:0
00 10 * * * /usr/bin/gedit
```

I recommend that you add the *DISPLAY* environment variable line on your *crontab* entry if you are scheduling to run applications that need X (GUI); else, it will not work.

Writing good *cron* jobs

The *cron* job line consists of the following five parameters separated by spaces:

1. Minute (from 0-59)
2. Hour (from 0 to 23)
3. Day of month (from 1 to 31)
4. Month (from 1 to 12)
5. Day of week (from 0 to 6) (0=Sunday)
6. Command

In most cases, you may find the asterisk (*) given as the field. If the hour field is given as * and the minute field as 0, it means that the command will run every hour at 0 minutes, i.e., * means every minute, or hour, or day of month, or month, or day of week.

Take a look at the following sample *cron* job entry:

```
00 02 * * * /usr/bin/ktorrent
```

What does the entry mean?

- 00 – 0th minute
- 02 – 2am
- * – any day of month

- * – any month
- * – any weekday

The above *cron* job can be translated to structured English as follows:

Execute `/usr/bin/ktorrent` at the 0th minute of 2 a.m., on any day of the month, any month, any day of the week.

The *cron* syntax allows you to specify the parameter for each field with hyphens to specify the range, i.e., for minute files you can specify 0-10. It also permits you to use comma separators to specify multiple parameters for the same field—i.e., 0,3,6 0-5 *** command is a valid *cron* job. The syntax also allows you to add comments in the *crontab* entry.

```
# Open gedit at 10 am
00 10 * * * gedit /home/slynux/sample.txt
```

I'm sure you now have some idea on how to write *cron* jobs. Hope you make the best of the unlimited bandwidth that's only made available at certain times of the day.

More *crontab* recipes

To start downloads at 2 a.m. and stop them at 8 a.m., execute the following:

```
00 02 * * * cd /home/slynux/distros/; wget -c http://
example.com/ubuntu.iso
00 08 * * * killall wget -s 9
```

To execute *diskusage.sh* every 30th minute, repeatedly, use the following command:

```
00/30 * * * * /usr/bin/diskusage.sh
```

To update the *locate* command search database every Sunday between 8 a.m. and 8 p.m., add the *crontab* as the root user, since *updatedb* needs higher write privileges:

```
* 8-20 * * 0 updatedb
```

To schedule system updates at 12 a.m. (here, too,

set *crontab* as the root for obvious reasons), use the following command:

```
00 0 * * * apt-get dist-upgrade -y
```

...if you use a Debian-based distro.

Or:

```
00 0 * * * yum upgrade -y
```

...if you have a Fedora-based one.

To shut down your machine at 10 p.m. (install *crontab* as the root user):

```
00 22 * * * halt
```


To remove *crontab* for a specific user:

```
# crontab -u username -r
```

Logging outputs from commands that you run

You can record the progress of the commands run by redirecting the standard output to some log file. It will be very useful to trace or debug something unusual that interrupted the *cron* job. Write your *cron* job as follows:

```
00 00 * * * command >> /var/log/cronjob.log
# cat /var/log/cronjob.log # For viewing log file.
```

cron is a utility that gives you an awesome user experience. There are also GUI implementations for it. But it is always fun to do everything on the command line, as it powers you to unleash the ultimate potential of the GNU/Linux system. There is also a utility called *at* for temporary job scheduling. Have a look at *man at*. That's it for now. Have fun with *crontab*, and happy hacking! 

By: Sarath Lakshman is an 18 year old hacker and free software enthusiast from Kerala. He loves working on the awesome GNU/Linux environment and he contributes to the PiTiVi video editor project. He is also the developer of SLYNIX, a distro for newbies. He is currently studying at Model Engineering College, Cochin. He blogs at www.sarathlakshman.info





INDUSTRY NEWS

Open source enters the mainstream according to a survey

Actuate Corporation, a Rich Internet Applications provider, has revealed some of the findings of its Actuate Annual Open Source Survey for 2008 conducted across the UK, North America, Germany and France. The survey, now in its third year, provides a global benchmark of attitudes and trends in open source growth and adoption.

The 2008 Actuate Annual Open Source Survey confirms that Europe in particular is forging ahead in widespread adoption of open source software, having recognised early on the lower cost of ownership, and the flexibility it offers for future application expansion and development.

Exploring in depth organisations' use of and attitudes towards open source, across four important territories, the findings categorically confirm that open source software is not a 'here today, gone tomorrow' phenomenon; rather, it has been broadly recognised and embraced for its ability to offer organisations sustained competitive advantage. These findings support Gartner's projections that, by 2012, at least 80 per cent of all commercial software solutions will include substantive open source components.

Close to 1,000 business and IT professionals took part in the survey, which was independently conducted by Survey Interactive in June 2008, and there were significant responses from financial services and public sector markets in the UK, North America and Germany. France and Germany also surveyed the manufacturing sectors in their respective countries. The survey concentrated on three key areas within open source software—awareness and adoption levels; benefits and barriers to adoption; and the evolution of open source Business Intelligence. Some highlights of the survey are that Europe leads in preference for open source platforms, particularly in the deployment of new applications, and replacement of outdated systems, with France and Germany at the forefront. Close to two-thirds (61.6 per cent) of French respondents said that open source is the preferred option while procuring software. This is a significant statistic, exceeded only by Germany (63.6 per cent).



MySQL co-founder quits Sun



David Axmark, co-founder of MySQL AB, the company behind the MySQL database, has quit Sun Microsystems. MySQL AB was acquired by Sun Microsystems for \$1 billion earlier this year.

In his letter of resignation, he wrote: "I have thought about my role at Sun and decided that I am better off in smaller organisations. I HATE all the rules that I need to follow, and I also HATE breaking them. It would be far better for me to 'retire' from employment and work with MySQL and Sun on a less formal basis."

According to a blog entry by Kaj Arno, VP of community for MySQL AB, Axmark will be "...working as a consultant for Sun, doing speaking engagements and connecting us with his huge network."

FSF 'reboots' its high priority list

The Free Software Foundation (FSF) has announced a 'reboot' of its High Priority Projects list with an accompanying \$10,000 grant from Worldlabel.com Inc. The grant will seed a new fund to promote projects on the list, and the FSF is calling for a community conversation about the biggest challenges computer users face with free software.

Russell Ossendryver, owner of www.WorldLabel.com, said, "Smaller companies and individuals can pool their resources in support of critical free software projects, but awareness is key. There are many threats from proprietary software and I wanted to contribute to a program that can help solve those problems. I am looking forward to working with the FSF to find creative ways to promote the cause."

FSF campaigns manager Joshua Gay, emphasised that the list is not considered static or complete, and that the FSF is seeking community input. "The FSF is asking the community of free software users who understand the critical issues that free software faces to tell us about the areas where they face problems. Problems that affect the most users are of the highest priority."

The list is online at <http://www.fsf.org/campaigns/priority.html>. It includes Gnash, a project to replace Adobe's proprietary Flash player; Coreboot, a free software replacement for proprietary BIOSs; a call for a free software replacement for the VoIP and multimedia chat program Skype; a free software membership and donor transaction and contact system for non-profit organisations; a free software replacement for Google Earth; and several more.



INDUSTRY NEWS



The CME Group joins Linux Foundation

The CME Group, the world's largest and diverse derivatives exchange, has joined the Linux Foundation—a non-profit organisation dedicated to accelerating the growth of Linux. The CME Group has been recognised as one of the financial services industry's biggest users of Linux. It first started using Linux in 2003 in order to cut costs, increase reliability, and reduce the round-trip time of a trade transaction. Since then, broader use and newer versions of Linux, coupled with match engine and application improvements have helped continue that trend. In an industry where low latency is paramount, this reduction extended the fundamental savings of Linux by enabling more transactions to be made in a given day.

As per the CME Group, last year it traded a record 2.2 billion contracts worth more than \$1.2 quadrillion.

By joining the Linux Foundation, the CME Group will be able to collaborate with key Linux developers and vendors. CME Group's Vinod Kutty, associate director and head of distributed computing R&D, will become chair of the Linux Foundation's End User Council. The Council is a group of the largest Linux end users who use the forum to collaborate and educate themselves on technical, legal and community efforts.



Novell to contribute to Moblin

Novell is going to participate in and contribute to the Moblin open source project. The company will start immediately contributing to Moblin's specific sub-projects, such as its desktop compositing manager and multimedia abstraction layer. As a result of Novell's participation, Moblin will be able to leverage Novell's large community of contributors and many existing open source projects. Novell will promote Moblin within these projects and continue its role in ensuring future innovations for all Linux desktops, including mobile and connected devices.

Moblin.org is an open source community for sharing software technologies to create an untethered computing experience across Mobile Internet Devices (MIDs), netbooks, and embedded devices using computing hardware based on Intel Atom Processor technology. Novell's focus is on effectively meeting the requirements of hardware vendors choosing the Intel Atom processor technology for their netbooks to deliver solutions with greatly accelerated boot performance, longer battery life and an optimised experience for netbook users.

"Novell's Moblin-based products will continue their track record of innovation," says Doug Fisher, vice president, Software and Services Group, Intel Corporation. "Delivering Moblin solutions on Intel Atom processor based platforms will take advantage of the new opportunities with rich Internet experiences designed for users on the go."

Qt, TietoEnator enter partnership

TietoEnator, a professional service company providing IT, R&D and consulting services, and Qt Software (formerly Trolltech) have signed an agreement for TietoEnator to use the Qt application framework in TietoEnator's Telecom and media development centres. TietoEnator Telecom and Media will establish competence centres for Qt in the Nordics, EMEA and Asia-Pacific that will offer professional services to customers worldwide.

"We anticipate good growth in the adoption of Qt for software development projects, especially in light of Nokia's acquisition of Trolltech," said Ville Aittomäki, senior vice president, Telecom & Media, Mobile Devices R&D of TietoEnator. "We are already engaged with several customer projects involving Qt, and believe that the demand for this competence will increase, going forward."

Serena acquires OSS alternative to MS Project

Serena Software has acquired Projity for OpenProj, its free, open source desktop project management software, and Project-ON-Demand, its Software-as-a-Service (SaaS)-based project management software. With this acquisition, Serena aims to bolster its Serena Mariner Project and Portfolio Management offering.

Rated as a leader by Forrester in project portfolio management tools last year, Serena Mariner provides total visibility into project and portfolio status and metrics to ensure the right people are on the right projects at the right time, delivering more value to the business. With the addition of the Projity offerings, Serena Mariner gains an alternative to Microsoft Project.

INDUSTRY NEWS

Mandrive, Turbolinux to contribute to Moblin

Mandrive and Turbolinux have announced the creation of Manbo Labs, which will join the Moblin.org project to contribute their technologies. Moblin.org is an open source community project to develop the Linux software stack and technology framework for visually rich Internet media experiences on devices such as Intel Atom Processor-based Mobile Internet Devices (MID), netbooks and in-vehicle infotainment systems.

The two will develop the core components of a Linux distribution based on Moblin.org. Both already have unique technology and strong expertise in developing netbook-oriented Linux distributions that will contribute to Moblin.org. The new Manbo Labs distro for netbooks will be Moblin-compliant and thus will be optimised for the Intel Atom processor platforms. The features that it will include are power management enabling long battery life, fast boot time, a small footprint and a slick GUI interface adapted to the small screen.

Hughes selects Wind River hypervisor

Hughes Telematics has selected Wind River to deliver the foundation for its next-generation telematics architecture. Specifically, Hughes' telematics control unit (TCU) will deliver information to drivers using Wind River Linux, a commercial-grade Linux operating system, and the recently introduced Wind River hypervisor will enable hardware virtualisation. Additionally, Wind River will provide professional services to the Hughes software development team throughout the development process.

Greg K-H attacks Canonical

Greg Kroah-Hartman, a core Linux kernel contributor, has shot off many salvos at Canonical, the company behind one of the most popular GNU/Linux distributions during The Linux Plumbers Conference.

He started by saying that the contributions of Canonical are minuscule compared to other players. According to his own blog, "A few months ago I gave a talk at Google about the Linux kernel development process. During that talk, someone asked me about Canonical's kernel contributions as they did not show up on the list that I was showing. I offhandedly remarked that they did not show up as they had only contributed 5-6 patches in the past few years. Now this comment didn't go over very well with the Ubuntu developers, and they called me out on it as they felt it was wrong."

He added, "In the past three years, from the 2.6.15 kernel to 2.6.27-rc6, Canonical has had 100 patches in the Linux kernel. I apologise about my previous statement and would like the world to know the correct number here. But as the Canonical employees seemed so eager for me to get the number correct, let's look a bit closer at it. What does 100 patches really mean?"

Matt Zimmerman, who currently works for Canonical as the technical leader of the Ubuntu project, and chairman of the Ubuntu technical board and CTO of the project, has tried to make the picture clear. He wrote on his blog, "Greg considers the 'Linux ecosystem' to be GCC, binutils, the Linux kernel, X.org, and a handful of other projects. He disregards most of the desktop stack (including GNOME and KDE), all desktop and server applications, and most anything else that is recognisable to an end user as 'Linux'."

He wrote on his blog that, "However, no one, certainly not Canonical, has ever claimed that Canonical does as much Linux development as Red Hat or Novell. He's refuting a claim that has, quite simply, never been made. Canonical is primarily a consumer of the Linux kernel. It is one of the building blocks we need in order to fulfil our primary mission, which is to provide an operating system that end users want to use. It is, on the whole, a good piece of software that meets our needs well. We routinely backport patches from newer kernels, and fix bugs which are particularly relevant to us, but our kernel consists almost entirely of code we receive from upstream."

It seems that Canonical has taken the criticisms about Ubuntu's lack of contribution upstream quite seriously and announced the beta of the Ubuntu Upstream Report. "The upstream report is a real-time page that lists the Top 100 projects in Ubuntu sorted by open bugs, and also shows us how many of those bugs are 'upstreamable', and how many of those have an actual link to an upstream bug tracker," according to a blog by Ubuntu's Jorge O. Castro. "The intent of this page is to be used as a tool by Ubuntu developers to track how well their linkages to upstream bug trackers are, and provide a list of bugs that are possible 'targets' for them to push upstream."





Download Them All With 'Wget'



Want to download an ISO, some music files and a few huge movies? Or mirror an entire website and just some pages off another site? Wget will take care of all your downloading needs.

Wget [www.gnu.org/software/wget] or GNU Wget (derived from *World Wide Web* and *GET*) is a network utility to download files from servers, and mirror websites using *http*, *https* and *ftp* protocols. The main reason you should use Wget is that it's very simple. From the command line or via the GUI, it packs a punch!

Features

Some of the compelling features of Wget are:

- Can resume aborted downloads
- Can use filename wild cards and recursively mirror directories
- Supports HTTP proxies and cookies
- Supports persistent HTTP connections
- Uses local file timestamps to determine whether documents need to be re-downloaded when mirroring

You can access the detailed listing of features by running *man wget* at a terminal prompt.

Getting Wget

If you use Linux, (Open) Solaris or any of the common UNIX flavours, chances are that you already have 'Wget' installed. Mac OSX and Windows users will have to compile the source to start using Wget.

The Wget source code is available at ftp.gnu.org/gnu/wget for download.

Using Wget

A generic Wget command line looks like:

```
wget [option]... [URL]..
```

Let us now look through a few 'Wget' one-liners:

1. Downloading a page/file/application:

```
$ wget http://www.http://www.pendrivelinux.com/  
$ wget ftp://ftp.gnu.org/gnu/wget/wget-1.11.2.tar.gz  
$ wget https://addons.mozilla.org/en-US/
```



```
thunderbird/downloads/file/26261/lightning-0.8-tb-linux.xpi
```

2. Wget supports resuming the download of a large file like an ISO with the `--continue` (`-c`) option. Resume with the same switch.

```
$ wget -c http://nginyang.uvt.nl/hardy/ubuntu-8.04-dvd-i386.iso
```

3. Mirrors websites: If you simply want it all, use the `--recursive` (`-r`) switch, but make sure you have enough disk space!

```
wget -r http://en.wikipedia.org/wiki/Main_Page
```

To mirror a website with links up to, say, three levels, use the `--level` (`-l`) option. The default maximum depth is five:

```
wget -r -l 3 http://en.wikipedia.org/wiki/Main_Page
```

To download all the pages (`-r`, recursive) on Wikipedia, plus one level (`--level=1`), into any other sites it links to, use the `-H` (span hosts) option, and to convert the links in the downloaded version to point to the other sites in the downloaded version, use the `-k` option. To get all the components like JPEG images present in each page, use the `-p` option. The following is the aggregate of this example:

```
$ wget -H -r --level=1 -k -p http://en.wikipedia.org/wiki/Main_Page
```

4. Some websites ask for the username and password to download a file. Thankfully, Wget supports this option:

```
$ wget --http-user=your_username --http-password=your_password url
```

The only security downside is that your password is visible. This is not a serious problem for home users, though.

5. Downloading multiple files is also possible with Wget. Simply use the URL within single quotation marks:

```
$ get 'url1' 'url2' 'url3'
```

6. Some websites may be unresponsive, so use the `-t` option to resolve this. To try 3 times use:

```
$ wget -t 3 url
```

7. When interacting with the network, Wget can check for timeout and abort the operation if it takes too long. Use the `--timeout=seconds` (or `-T`) option. Here, we have used a timeout of six seconds.

```
$ wget --timeout=6 url
```

8. Time stamping: If you would like to download a file so

that it keeps its date of modification, follow these steps:

```
$ wget -S http://www.gnu.ai.mit.edu/
```

`ls -l` shows that the time stamp on the local file equals the state of the last-modified header, as returned by the server.

Later, you would like Wget to check if the remote file has changed, and download it if it has changed; use the `--timestamping` (`-N`) option:

```
$ wget -N http://www.gnu.ai.mit.edu/
```

9. If you have a file that contains the URLs you want to download, then use the `--input-file=file` (or `-i`) switch. If this function is used, no URLs need be present on the command line:

```
$ wget -i filename
```

10. To get Wget to download all JPEG images at a site, use the `--accept=file_type` (`-A`) option as:

```
$ wget -r --accept=jpg,jpeg www.sitenam.com
```

11. To recursively download a file and convert links in HTML files to point to local files for offline browsing, use the `--convert-links` option:

```
$ wget --convert-links -r http://www.mypclinuxos.com
```

There is much more Wget can do, like using proxies, using certificates to check before downloading, etc, making it one of the most versatile download managers.

Beyond CLI

Gwget [www.gnome.org/projects/gwget] is a free graphical download manager for GNOME, which uses Wget as its backend.

Doing more with Wget

Check out the Wget manual page (`man wget`), or www.gnu.org/software/wget/manual/wget.html and `wget --help` from the command line to see all that it can do.



MORE INFORMATION:

- Wget Homepage: www.gnu.org/software/wget
- Wget FAQ: wget.addictivecode.org/FrequentlyAskedQuestions?action=show&redirect=FAQ
- Wikipedia entry on Wget: en.wikipedia.org/wiki/Wget#Features

By: Nelson Lobo. The author is a lecturer at St.Xavier's College, Mapusa, Goa. He can be reached at lobonanc@gmail.com

Parsix GNU/Linux A Gift From Persia

Hey, all you Ubuntu folks! Open your eyes. Strain your ears. Listen. Ubuntu just lost the war.

We all know about the kitchen-sink Linux distro—Knoppix. It includes a lot of software and that becomes a disadvantage, because the software we don't use, accumulates as bloat. Useful though it may be, Knoppix is enormously unwieldy.

To address this problem, Kanotix was created. The Kanotix team took Knoppix apart, removed more than 80 per cent of the software, and repackaged it in a home-user friendly package. Kanotix was a jewel at one time. It gave us the power of Knoppix, yet made it usable for the novices.

But Kanotix died an untimely death. Lack of community support and the lack of time for its author were the main causes. But the structure of Kanotix was so rigid and robust that some guys from Iran got together and took up the mammoth task of updating Kanotix regularly. Finally, they gave a new name to this distro—Parsix.

The LiveCD

Parsix comes on a single CD. This is quite surprising, considering the amount of data in it. Pop in the CD and reboot, and a menu comes up asking you to start or install Parsix. It also has options for two widescreen resolutions, checking MD5 hashes for the CD and starting installation in text mode. Anyway, I just hit *Enter* and watched the CD boot up. On my modest 2.4GHz Pentium 4 'Prescott' with 512 MB of RAM, the CD took an impressive 32 seconds to start X. From then on, it took another 7 seconds to finish booting





Figure 1: Default desktop

the system and keep GNOME running. This is without a swap partition.

Then came the first problem—NTFS drives. Parsix Live boots as a normal user Parsix, and NTFS-3G needs super-user privileges to access NTFS drives. Blast!

Then I started looking at the software that comes bundled with this distro. Table 1 sums it all up. Please note the Iceweasel browser, which is just the same as Firefox, at version 3.0.1 and the Gnash SWF Player. Gnash was a bit of a surprise to me, considering that it was not stable enough. But it's good.

Installation

Now here comes Parsix's biggest flaw: the installation procedure can confuse even the most experienced of Linux users. I almost succeeded in corrupting my entire partition table trying to just format my existing openSUSE 11 drive.

For a successful installation, follow these few not-so-simple steps:

1. Download and burn the Parted Magic CD Image (see *Resources*).
2. Boot up your system using Parted Magic and create/edit/format your Linux and swap partitions.
3. Now boot up your system using the Parsix LiveCD and click on the *Install Parsix* icon on your desktop. If you can't start the installer, open a terminal, and type the following:

```
su -
parsix-installer
```

4. Click *OK* to get to the main menu.
5. Choose *Configure Installation* and click on *OK*.
6. Choose a root partition. No, you cannot use a separate home partition. Click *Next*.
7. Now choose a filesystem. I'd have loved XFS, but had to settle for ReiserFS. Other choices are ext3 and JFS. Click *Next*.
8. Now enter your name in the format (name, followed by surname), as in 'Boudhayan Gupta', and click *Next*.
9. Then enter your username and click *Next*.
10. In the next two screens, enter and confirm your user

TABLE 1: LIST OF APPLICATIONS

Categories	Applications
Multimedia	VLC Media Player, Exaile, Gnash, Sound Juicer, a certain TV viewer software that refused to start, GNOME CD player
System	Standard GNOME Utilities Package, GParted, VirtualBox OSE, Compiz-Fusion Icon
Games	Standard GNOME games package
Office	Openoffice.org Writer, Calc, Impress and Draw, Grisbi, eFax GUI
Graphics	Cheese, The GIMP, gThumb, XSane, Inkscape, Evince PDF Reader
Internet	Iceweasel Browser 3.01, Balsa E-Mail Client, Firestarter Internet Firewall, XChat IRC Client, Transmission BitTorrent Client, GWget Download Manager, GNOME FTP Client, Terminal Server Client, Pidgin IM, LifeRea FeedReader
Parsix Configuration Tools	CUPS, Documentation, Network Configuration, PPP Configuration
CLI Application Versions	Coreutils 6.10, Binutils 2.18.0, GNU Compiler Collection 4.3.1, Kernel 2.6.24 with SMP support and Pre-Emptive Multi-tasking Support

and root passwords.

11. Now choose where the bootloader should be installed.

Parsix can peacefully coexist with all major Linux distros and all versions of Windows. They will be auto-detected when GRUB 0.97 is installed.

Now choose *Start Installation* from the main menu and confirm your installation configuration. Then sit back while Parsix is installed.

Leave the other options (especially partitioning) alone (see side box).

Further exploration

So after installation, I wandered around the system trying to find out its nooks and crannies. I came up with quite a few.

First of all, I started by changing the wallpaper. To my surprise, only the default Parsix wallpaper and the Debian Swirl wallpaper were available.

Then came the themes. Again, only Debian, Parsix Viola (v1.5, this one), Parsix Ramon (v1.0) and Ubuntu Human themes were available.

There are quite a few interesting applications, including the full version of VirtualBox 1.6.2. But try Xnest. You can immediately log in as another user in a new Window.

Compiz Fusion is also included, but it's so tucked away that it took me ages to find. Also, Emerald needs more themes.

Multimedia



Figure 2: Multimedia applications




Figure 3: System tools

This needs a section devoted to it. First of all, it plays MP3s at full quality. Then, there is more to this section than meets the eye.

Since Amarok is missing (no KDE), I had to settle with Exaile. I was impressed with the sound quality, especially for the MP3s. I played MP3s of Linkin Park's latest album *Minutes To Midnight*, to test it. MP3s on Linux have a tendency to make the speakers sound like they are exploding. MP3s on openSUSE 11 sounded terrible on my system. I also amplified any possible effects by turning both the PCM and Master values on the Mixer to 100. To my surprise, no exploding effect occurred.

Instead of the default Totem, Parsix comes with VLC Media Player. Another pleasant surprise, since VLC has a reputation of playing anything that is thrown its way. It successfully played the FLV of the Large Hadron Collider that I had downloaded from CERN's website. It also played MP4s downloaded from YouTube (H.26x). It includes a TV viewer, but MythTV would have been a better choice.

I tested unprotected WMAs with both Exaile and VLC. Exaile did not open them, and VLC played them without a glitch. But there was the speaker explosion effect that put me off. But if you keep the VLC Volume to 29 per cent or so and turn up the speaker volume, the explosions are minimised.

And yes, iRobot from a DVD in VLC looks and sounds cool, so does Nancy Drew in full High-Def format. I did not test protected DVDs, but it should work as *libdvdcss* is

WHAT HAPPENED WHEN I GOT 'PARTITION-HAPPY'

When I installed the distro, I went to the partition menu and started partitioning for my drive, which is sda. Nothing happened, so I started GParted manually (GParted is supposed to start). Then I unmounted my openSUSE XFS partition and set it to format using ReiserFS. I clicked on Apply, and everything seemed to go along happily, until mkfs failed and GParted quit with an error message about partitioning having failed. I repeated the procedure multiple times, with the same result. I found out that just as I had unmounted the Linux partition, HAL had remounted it again. So I ditched GParted and went the old cfdisk way, preferring to remove and recreate my Linux partition, which I did. And I successfully corrupted my entire partition table. AAAAAAAAAARRRRGGGGHHHHH!

I then went back to my trusty old Parted Magic 1.9 (an antique, currently at 3.0) CD and used testdisk to recover my system MBR. Relief!

included by default.

The big problem lurking under the hood

No, it's not a problem. It's the kernel, actually—an old 2.6.24, modified on the already modified Ubuntu sources. An absent Bootsplash was another thing that disappointed me. If you are up to it, I'd strongly recommend that you get the latest kernel sources from kernel.org, patch it with Bootsplash (not necessary) and compile your own kernel. If you can't, that's not a major problem either!

More 'to-dos'

There are a few things you'd most probably want to do before you start using the system. First of all you might want to include the Lenny repositories and run an *apt-get dist-upgrade*. This will update your system to the latest packages in the Lenny repositories. You can also include Ubuntu, Kubuntu, Xubuntu or Ubuntu Studio as they are all binary compatible (Google-search for Linux Standard Base) but the Ubuntu, Xubuntu and Kubuntu repositories would be useless as the Lenny repositories already have the required software in later versions.

Then you might want to configure Compiz-Fusion and Emerald. They are not enabled by default, and they need manual configuration. Head to *Applications-->System Tools-->Compiz Fusion Icon*. Then right-click on the Compiz icon on the panel and adjust the settings in the settings manager and install an Emerald theme. Once you are done, select the window decorator to be Emerald and the window manager also to be Emerald. Then reload the window manager and you are done. Now set Emerald as the default window manager to retain the settings.

You need to install the ATI or NVidia drivers manually, if you have those chipsets, because they are not included by default.



Performance

Ah, here comes the big test. And Parsix passes it beautifully. I have quite a powerful processor, a 2.4 Ghz Pentium 4 'Prescott', which happens to support SSE3, but everything else about it is crappy. Just 512 MB of RAM and a rotting old Intel 865 chipset with integrated 82865 G Graphics make it an average system. Above that, I have no swap partitions and no swap file.

And Parsix works faster than DOS here. To tell you the truth, I've tuned Compiz to use the most processor heavy effects and I'm using a very compositor-happy theme, and the system is as responsive and as smooth as anything. With only 55 per cent memory usage at full load and negligible processor usage, it's racing faster than an F1 car.

Nothing has even remotely slowed down or crashed, as yet, and that's after almost 10 days of use. This is beginning to be too good to be true. Maybe I'm overreacting after my particularly nasty experience with openSUSE 11, but Parsix is a hell of a performer.


The judgement day

The following scores (out of 10) are what I give to Parsix (compared to Ubuntu 8.04 LTS):

1. Performance: 9.5 vs. 9
2. Reliability: 10 vs. 8
3. Usability: 7 vs. 4
4. Looks: 8.75 vs. 6.5

5. Included Software: 7.75 vs. 7.5
6. Configuration Tools: 5 vs. 5
7. Multimedia: 7.75 vs. 3
8. Overall: 9 vs. 7

The final judgement

The Iranians are good with FOSS stuff. 

RESOURCES AND REFERENCES:

- Parsix: www.parsix.org
- 32 Bit ISO: nchc.dl.sourceforge.net/sourceforge/xfardic/parsix_1.5r1-i386.iso
- 64 Bit ISO: nchc.dl.sourceforge.net/sourceforge/xfardic/parsix_1.5r1-amd64.iso
- Parted Magic: www.partedmagic.com
- 32 Bit ISO: partedmagic.com/downloads/stable/pmagic-3.0.iso.zip
- Debian: www.debian.org

Some Screenshots are from Parsix's Site and The Coding Studio.

By: Boudhayan Gupta is a 14-year-old student studying in Class 8. He is a logician (as opposed to magician), a great supporter of Free Software and loves hacking Linux. Other than that, he is an experienced programmer in BASIC and can also program in C++, Python and Assembly (NASM Syntax).

LINUX

ForYou

Calling...

Share your knowledge

Inviting FOSS experts to write articles on their area of interest

LFY covers a myriad of topics—network management, software development, embedded systems, community issues, and even hands-on guide for newbies. If you've got an interesting topic, let us know. Thanks to the launch of www.openITis.com (aka linuxforu.com) we are now trying to extend our content portfolio related to Linux & Open Source.

To know more on how to become an LFY author, contact us at
lfyedit@efyindia.com

Few topics that top our list:

- ◆ Tips 'n' Tricks for software developers or IT implementers
- ◆ Cool tweaks for FOSS enthusiasts
- ◆ FOSS on mobile
- ◆ Virtualisation (Implementation)
- ◆ OpenJDK or Java on Linux
- ◆ OpenSolaris (software development)
- ◆ How can I do 'that' on Linux
- ◆ Reviews of latest open source projects & tools





Jaijit Bhattacharya, country director, public policy and corporate strategy, Sun Microsystems India

The ODF Olympiad

The End of the Opium War?

The ODF Olympiad is back again. We caught up with Jaijit Bhattacharya, country director, public policy and corporate strategy, Sun Microsystems India, to not only get details about the contest, but also understand why the Open Document Format (ODF) matters so much for the next generation.

Q Why is the Open Document Format so important for the next generation?

The ODF is the corner stone of technological sovereignty. It is also a tool to reduce the digital divide and a path towards the adoption of next generation open source technologies.

If we look at the manner in which technology is being controlled by a few firms (which are primarily based in the West), it is evident that firms use the vicious cycle of getting undue profits from monopolistic standards (which can be *de facto* or *de jure*). These are then pumped into newer formats and standards that are sometimes pushed through international standards bodies such as the ISO, which are then pushed down the throat of consumers in both developed and emerging economies such as India. And then the cycle continues. Ever so often, such companies unilaterally declare their own

formats obsolete and push in newer formats that force people to upgrade PCs, thus also leading to the generation of enormous volumes of e-waste.



If one adopts ODF, one can start using cutting-edge free and open source software such as OpenOffice.org, which is not only technologically more advanced, but is also free of cost. This reduces the cost of accessing computer education and hence reduces the digital divide.

And finally, if you look at the largest Internet firms such as Google, YouTube, etc, you will notice that they are being powered by open source software such as MySQL. Hence, if we are denying education in open source software to our children, we are denying them the skills to create the next Googles, YouTubes, etc.

Q: Among what profile of people is it most important to promote ODF?

From the earlier discussion, it is quite clear that the vicious cycle of proprietary software and disproportionate economic benefits to a few corporations, starts from the schools. As the Chinese government refers to the phenomenon, it is the next 'Opium War'. Our children are being given the opium of 'proprietary software' from their school days and are being tested on their proficiency of using proprietary software, so that when they grow up, they are capable of primarily using proprietary software. The situation is even worse when one looks at the fact that the monthly income of a majority of people in the country is less than Rs 3,000, and they are forced to pay for proprietary software whose cost is more than their monthly salary.

Therefore, it is an urgent need to ensure that our children are free from proprietary software, and are given an education using free and open source software. This will also enable them to set up the next generation of Internet firms.

Q: What role can schools play in further promoting ODF?

As I said, the school is the starting point in the vicious cycle of pushing proprietary software into the economy and extracting undue economic benefit by manipulating the working environment and the standards.

When we conducted the first ODF Olympiad, schools came back and asked why they are prevented from using OpenOffice.org as a teaching medium for word processing. We soon discovered that even the question papers test skills in proprietary software and not in open source. It is necessary and essential for schools to quickly switch over to free and open source software immediately in order to give the children a competitive edge.

Q: What is the ODF Olympiad and how can schools ensure participation?

The ODF Olympiad is a multi-stakeholder initiative being promoted by the government of Malaysia, the government of West Bengal, Sun Microsystems, GNU, the Free Software Foundation, the ODF Alliance, IBM, Red Hat, IIT Delhi, JNU, etc.

FACTS FOR YOU

Topics and categories

- Up to Class V -Why should I learn computers?
- Class VI to VIII -Benefits of Open Source Software to my school
- Class IX & X -Open Standards: Freedom from the Digital Divide
- Class XI & XII -Importance of Open Standards for good governance

Evaluation parameters for the ODF Olympiad trophy

- Relevance of topic
- Creativity and visual appeal
- Organisation of the content
- Use of embedded tools (e.g., spreadsheet)

ODF Olympiad process


- Create presentations in ODF using OpenOffice.org, or any other free/libre and open source ODF-compliant software.
- The presentation should not have more than a total of 21 slides (including the introduction and 'thank you' slide).
- The presentation file size should not exceed 1 MB.
- Each slide should not have more than two clicks for navigation.
- E-mail the presentation to india@odfolympiad.org along with your personal details—name, class, school's name, school's address, e-mail address, phone number and age—before November 15, 2008.
- The subject of the e-mail must be in the following format: category of presentation, school name, for example "Category II, Gandhi National School".

The ODF Olympiad aims to showcase the ease of using cutting-edge ODF-compliant software such as OpenOffice.org and thus bring students and teachers closer to adopting free and open source software. We had an overwhelming response the first time we conducted the ODF Olympiad and we were flooded with requests to continue having the ODF Olympiad every year.

It is very simple to participate. The instructions are available online at www.odfolympiad.org. All that students in India have to do is submit a 21-slide presentation on one of the four topics.

Q: What role is Sun playing in this endeavour?

Sun is hosting the secretariat of the ODF Olympiad and is part-sponsoring it. We are also helping in spreading the word around, and in coordinating with the various agencies and stakeholders involved.

For clarifications, e-mail to secretariat@odfolympiad.org 

By: Swapnil Bhartiya, assistant editor, EFTimes.com



Dare To Let That **Puppy** Loose?



In this calorie-burning, fat-avoiding, health-conscious world, our poor computer gets loaded with gigabytes of fatty operating systems. Try to cut short the fat and you'll end up either with a feeble OS unfit for regular use, or with a fairly small, but difficult-to-move alternative. The ultimate solution of a portable, fat-free desktop is made by setting free the small, thin Puppy on your disk and letting it manage your system and serve you with all the needed desktop applications.

Puppy Linux is a minimalist distro that demands very little resources from users in order to unleash its power. It supports installation on any media, like CD, DVD, USB

Flash drive, internal and external hard drives—you name it. And to run Puppy happily, the hardware requirements are: Pentium 166MHz MMX Processor, and a 20x CD-ROM drive or a bootable USB drive, else boot floppy to boot from other devices—no hard disk is necessary!

Created by Barry Kauler, the Puppy was born on June 2003. Though it is small and thin, it stuns its users with its flexibility, usability and features. Puppy has its own personality, as it was made from scrap and has not been stripped down off any other OS.

Being small, it loads itself completely onto the RAM and runs from there. This makes for very fast program access times—almost nothing in my case. If Puppy cannot fit fully onto the RAM when booted off the CD/DVD, which might happen in computers with less than 128 MB of RAM, it loads in a swap partition (which you can create), or a swap file; else, it runs from media it was booted from (viz., a live CD). Puppy comes

in different sizes depending on its versions and makes (official and unofficial). The official release stays below 100 MB. The latest version, Puppy 4.1 'Dingo', released on October 6, 2008, is only 94.3 MB in size.

Puppy stands out from the pack of other distros with its unique multi-session live CD feature. You can save the sessions along with settings, documents and downloads in the very CD you booted from. Each new state is stored in separate sessions as directories. Sessions can also be saved in a USB drive or in a HDD with FAT or NTFS, where it gets stored as a single file with a Linux filesystem inside it. When booting, Puppy smartly searches all the accessible devices attached to the system, and automatically loads the latest session data and starts with the latest saved state.

What's inside?

You may not expect much from the 94.3 MB tiny Puppy, but it will astonish anyone with the variety of applications that cover all sections that fulfil most of our needs.

The desktop experience gets better with the simplicity and usability of the default,

BOOT PARAMETERS

When booting from the Live CD, the boot screen displays some boot parameters that are handy under special cases. They are all self-descriptive. You can control the loading of Puppy onto the RAM, blacklist unwanted saved CD sessions, start the command line only, and more. To get information on boot parameters, visit www.puppylinux.org/wikka/BootParms

light and fast Joe's Windows Manager (JWM) and the fully-functional ROX file manager. Other window managers like Fvwm95, IceWM, Xfce, Fluxbox, Enlightenment, and also KDE, can be installed. GNOME is still to be ported. If you want to give Puppy a personal touch, themes are available online, and can be selected from the Puppy menu.

Formerly, Puppy was based on GTK+1 and Tcl/Tk. Things have changed now—it supports all GTK+2 applications.

Starting with the documents section, a reduced edition of AbiWord takes care of formatted word editing, with Leafy and Geany managing plain texts. A PDF viewer (ePDFview), PDF converter (puppyPDF), and a Windows .chm help file reader are at your service. The need of a spreadsheet application is satisfied by the fully Microsoft-Excel-compatible Gnumeric. You will even find a personal finance application (HomeBank), Osmo personal organiser, personal wiki (DigiWiki), calculators, notes, address book, the Seymonkey Web page editor and more. And with CUPS pre-installed, your printers just await your configuring them.

As for graphics, Puppy gives you the ftoox image viewer, as well as mtPaint and a light edition of the well-known Inkscape for necessary image editing needs. Apart from these, it has the digital camera manager GTKam and the Xsane scanner manager.

Now let us come to one of the most sensitive sections in Linux desktops—multimedia. Puppy has an in-built Pmusic audio player and Gxine multimedia player, which make it an excellent out-of-the-box multimedia player with lots of codecs support, including MP3, Flash and encrypted DVD playback. Ripping audio/video? It isn't a problem at all with Pbcdripper, Pupdvdtol and ripperX tools. Editing and burning ISO images are just a few clicks away with Burniso2cd and isomaster. Pburn lets you compile and burn files onto CD and DVD, and even onto BlueRay discs.

Coming to the Internet section, the SeaMonkey Web browser and Ayttn (which replaced Pidgin from 4.1) solves the browsing and chat application needs. The Axel download accelerator and the Pwget downloader are ready for the mass download addicts amongst us. Additionally, they also have the Pctorrent and Pcreatetorrent for torrent management. The gFTP ftp client, the Slypheed e-mail manager, and even an excellent VoIP application (Psip), along with Puppy's own PPLOG Perl blog with built-in Hiawatha Web server, completes the Internet set.

One-click mounting and unmounting is driven by Pmount, and the excellent disk detection is thanks to the MUT utility.



Figure 1: Default Desktop

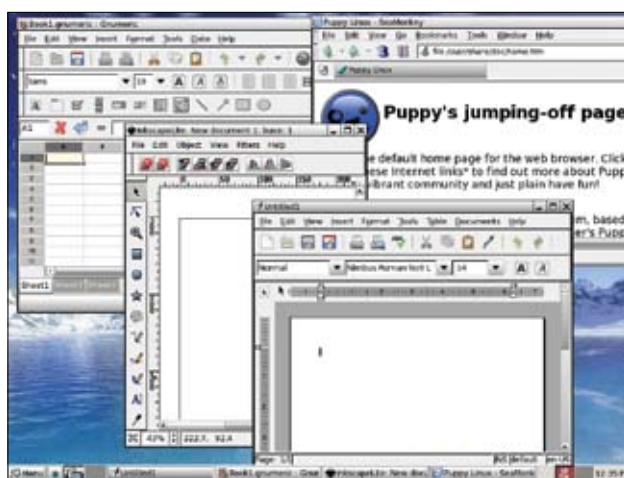


Figure 2: Running Gnumeric, Seaminkey, AbiWord and Inkscape on Puppy




Figure 3: Running Wizard-wizard, Gxine, ROX Filer and rvtx on Puppy

The Xarchiver compressor supports TAR, ACE, RAR, ZIP, 7ZIP, ARJ and many more popular formats. Among others, back-up utilities are batch file renamers (Prenam, gFnRename), a disk cataloguer (Gwhere), partition managers (Gparted, Pdisk), process managers (Pprocess), boot managers, terminal emulator (Rxvt, Xfproft virus checker), and a lot more to



SPECIAL CASES

You should be careful of whether the system suddenly shuts down while working with Puppy. Then you lose all your work as it evaporates from the RAM. So beware of power cuts.

When a CD/DVD is full, it just needs a new media to be inserted when writing. You can install Puppy in the new media before saving the latest session to avoid carrying the old CD for booting.

Puppy has problems loading X server with newer cards. I booted Puppy in my AMD Athlon64 X2 system with an ATI Radeon 1250X integrated graphics card. It booted into the console but failed to start the X server.

quench all your thirst. And don't forget to benchmark your system with the Hardinfo tool and test your RAM with Memtest.

Some of the assorted tools for your networking needs are PPPoE and dial-up connectors, Samba shares, Puppy Firewall. It also includes Rdesktop, which lets Puppy be used as a thin client with Windows XP, 2000 NT and 2003.

And, of course, as in all distros, Puppy also has some games in-built. Currently, it has a number puzzle, jigsaw puzzle and a Rubic's cube game. Obviously, more games are available on the Net. And would you believe me if I say that there are some more applications that you can discover that I have not included in the article?

Want more?

So, what good is a Puppy if you can't feed it. The developers might be diet freaks, but you are obviously allowed to make it fat by downloading applications available on the Internet using Puppy's own PETget package management system. PETget can get you more than 500 official and hundreds of unofficial applications. The best thing about Puppy is that it supports Debian and Slackware packages. With this support, Puppy's application range grows by a gigantic amount of more than 20,000 packages of Debian and Slackware.

Puppy's Squash filesystem support lets you download .sfs files of applications and load them directly at boot time, for use. You can download OpenOffice.org, JRE (Java Runtime Environment), and developer .sfs files from www.grafpup.org/download.php?path=.%2Fmodules%2F, which you need to mount after Puppy has booted.

What else?

Configuring and tweaking the system is like a breeze with 'wizard wizard', a very simple tool with a small window and some buttons that do not give you that awful feeling of, "Oh my god, how am I supposed to use this?" Keyboard, mouse, sound and graphics server, network, firewall, modem, CUPS, CD/DVD drives—which basically covers everything—can be configured with this simple interactive tool with just a few clicks. It really is the wizard's wizard!

A great feature of Puppy is that you can roll out a customised distro with just a few clicks on the Puppy

Remaster feature, which lets you create your own customised Puppy CD (or ISO image), with new packages, themes and all that you have downloaded. Thus you can quickly make your own Puppy release with applications of your choice like OpenOffice.org, Firefox, GIMP, and whatever you want.

The Puppy Universal Installer, which is in fact 'universal', offers you a variety of options to get Puppy installed in various types of media, graphically.

The best thing I liked about Puppy is the language it uses on the dialogue boxes, set-up options and everywhere. Very simple and every day language that makes a newcomer comfortable. Each message describes in detail and with great clarity, what an option is going to do.

And finally, unlike other Linux distros, it is a single-user system and you are always logged in as the root. Yet the community opinion is: this poses no security threat.

Ready to get started?

The download page www.puppylinux.org/downloads lists all the previous releases along with the current Puppy release. You will notice some 'puplet' releases as well. These are unofficial releases packed by Puppy users. Puplets are packed with specific themes, or are bundled with more applications. Check out the fatty 'Dingo Plus' Puplet containing OpenOffice.org, Firefox and more applications. Get lots of more puplets from www.puppylinux.org/wiki/archives/old-wikka-wikki/categoryderivatives/puplets

Anyway, we'll use the official version here. Download the Puppy 4.1 ISO file. You need to burn it in a CD/CD-RW. Make sure your burner software supports multi-session burning. Select the image to burn from your burner menu and make sure you enable the multi-session option before burning it. This lets you save Puppy files in the live CD. Similarly, you can burn a DVD if you need more space. For more information, please check puppylinux.org/wiki/how-tos/general/creatingmultisessionfrompuppy and puppylinux.org/wiki/how-tos/general/burnpuppycd

Now that you have the media ready, boot your PC. You are greeted by a boot screen with a list of boot parameters. Press *Enter* to boot! It searches for any saved Puppy session files in the accessible disks and copies them. Then it prompts for your keyboard layout—select the proper one (generally QWERTY). Then with the X server configuration, you have two choices: either the Xorg or the Xvesa. Xorg has better support for new hardware, but might not be the ideal choice for older machines. On the other hand, Xvesa is simple to configure, has limited functionality (doesn't have hardware acceleration), but you can expect it to run problem-free under any hardware.

In my case, I first chose Xorg in a test machine with Intel Pentium D 2.8 GHz, 512 MB RAM and an ATI Radeon 200x integrated graphics, which all worked fine, except that the Gxine video output was corrupt. Changing the X Server to Xvesa from the wizard-wizard solved it and let me enjoy Rocky Balboa.

After you have done this, the desktop loads within two seconds and you are greeted with a friendly "Woof-Woof". The

panel has the Puppy menu button, CPU and RAM monitor, and a run command button. You can switch between two virtual desktops. The desktop contains a set of organised application icons. From here, you can play around the desktop, for your own satisfaction.

To save the session after you have finished working, simply shut down the system from the Puppy menu. You will get options to save all the changes with *SAVE TO FILE*, which saves the session file in any disk drives accessible, or the *SAVE TO CD* to save the session in a CD/DVD. When using *SAVE TO FILE*, name the file, and make sure you select the correct place and the size of the session file (select as needed) when prompted. It also offers an option to encrypt the files to be saved. You need to keep the media connected with the session file that you want to be loaded when booting Puppy later on. When booted from the USB drive, there is a save icon on the desktop, which lets you save the session on the USB drive at any time.

Puppy Linux offers installation on the USB drive without formatting it or altering it in any way, while keeping all the existing data in place. But still, when installing it in the USB drive, I recommend you back up all the data, to be ready for any worst case that might occur. Launch the Puppy Universal Installer from the *Puppy Menu*→*Setup*. You are presented with a variety of choices about where you want to install Puppy. Select the USB Flash drive option and keep following the simple instructions. If you don't want to lose the existing USB data, then select 'No' when asked whether you want to format the USB drive. If your system can boot from the USB, don't create a boot floppy and continue with making the USB bootable.

Puppy HDD installation gives you two options: a Full installation that is a conventional Linux installation with a separate Linux partition(s) and a Linux filesystem, or the Frugal installation that lets you install Puppy in FAT or NTFS—even in the same partition with Windows, yet without any conflict. You need to manually install the bootloader later, though.

You can get more information at:

- puppylinux.org/wiki/how-tos/general/harddriveinstall
- puppylinux.org/wiki/archives/old-wikka-wikki/categorydocumentation/harddiskinstall
- puppylinux.org/wiki/how-tos/general/helpinstallingandbootingpuppy
- puppylinux.org/wiki/how-tos/general/hintandtips

Roll your own Puppy

After you have tweaked the settings, installed packages and customised Puppy, you can make a customised Puppy ISO by launching the *Puppy Menu*→*Setup*→*Remaster Puppy* live CD program by following some simple instructions. It lets you check and make final alterations to the final image. It even asks if you want to pack the hardware customisations as well.

A pre-hardware customised ISO might not work out-of-the-box on other hardware where you install this image; it's best not to save the hardware customisations. Saving hardware customisation helps if you are only going to use the distro with the same or similar hardware set-up.

Finally, you can either directly burn the customised image onto a CD or you might want to save it as an ISO file, which is saved in */inirtld/mnt/dev_save/custom-puppy-410.iso* (the number at the end may vary). You are now ready to distribute it!

The journey so far with Puppy has been a steady one. It's easy to use, comes with simple, everyday language, and its flexibility encourages users from all age groups. Its ability to be installed in any media makes it ultra portable. The multi-session live CD feature offers a cheap portability solution. The utilities section also makes it a very good system rescue CD. Except for desktop use, it also has features to be run on server systems. Puppy can stay in the same partition with Windows, making it the best way to make a Windows user try Linux. Most of all, it makes an old junk machine run again with its lost glory, while performing like a high-end computer. 

LINKS AND REFERENCES:

- Official site: www.puppylinux.org
- How to: puppylinux.org/wiki/how-tos
- Wiki: puppylinux.org/wiki
- Forum: www.murga-linux.com/puppy
- Downloads: www.grafpup.org
- Freenode IRC Channel: [#puppylinux](https://freenode.net/channels/#puppylinux)
- www.linux.com/feature/137880
- www.desktoplinux.com/articles/AT7455536044.html

By: Arjun Pakrashi. The author is currently studying for a B.Sc degree in computer science from Asutosh College, Calcutta University, Kolkata. His main areas of interest are open source software, Linux programming and data structures. He plans to do research-based work, and become an OSS contributor.

**Enterprise Mail Server, Linux SBS Server,
Anti SPAM, Antivirus and HTTP Filtering.**

**Bandwidth Management, Internet Access Control,
Content Filtering, Web Access Reporting.**

Customised Linux Product Development.

TechnoInfotech™

1, Vikas Permisses, 11 Bank Street, Fort, Mumbai, India - 400 001. Tel.: 91-22-6633 8900 Ext. 324. info@technoinfotech.com

CELEBRATE Diwali

Win EXCITING PRIZES

Subscribe Now!



Cover Price :
Rs 60/- with CD
Rs 40/- without CD*

A magazine for the electronics fraternity... professionals, businessmen and hobbyists

* From September 2008 issue



Cover Price :
Rs 100/-
with CD/DVDs

The complete magazine on open source



Cover Price : Rs 50/-
An IT magazine for businessmen



Cover Price : Rs 50/-
India's first Electronics Sourcing magazine



Cover Price : Rs 40/-
Information Technology—Today, Tomorrow.



Cover Price : Rs 100/-
India's most influential monthly on Economic and Business Affairs

Send the filled-in form (Opposite page) or its photocopy to :



EFY Enterprises Pvt Ltd, D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 020 Ph: 011-26810601-03;
Fax: 011-26817563, E-mail: subscribe@efyindia.com, website: www.efyindia.com

• Prizes will be selected through Lucky Draw. Gift Items photographs are indicative only. Condition Apply. • These rates are applicable for new subscribers as well as renewal by existing subscribers. • The rates are valid for subscribers within India only. • Please allow 4-6 weeks for processing of your subscription. • Disputes, if any, are subject to exclusive jurisdiction of competent courts and forums in Delhi/New Delhi only.

**1st
prize**

LCD TV



**2nd
prize**

**MP4 Player
[10 Nos]**



**3rd
prize**

**TRAVEL BAG
[100 Nos]**



**ASSURED
FREE GIFT**
with each subscription
**Round-Neck
T-Shirt**

Magazine	Unit Price (Rs.)	7 Years (Rs)	5 Years (Rs)	3 Years (Rs)	2 Years (Rs)	1 Years (Rs)
Electronics For You (EFY)	40/-	1680 <input type="checkbox"/>	1445 <input type="checkbox"/>	1015 <input type="checkbox"/>	725 <input type="checkbox"/>	385 <input type="checkbox"/>
EFY (with CD)	60/-	2520 <input type="checkbox"/>	2160 <input type="checkbox"/>	1510 <input type="checkbox"/>	1080 <input type="checkbox"/>	575 <input type="checkbox"/>
Electronics Bazaar	50/-	2100 <input type="checkbox"/>	---- <input type="checkbox"/>	1260 <input type="checkbox"/>	900 <input type="checkbox"/>	480 <input type="checkbox"/>
Linux For You (with CD & DVD)	100/-	4200 <input type="checkbox"/>	3600 <input type="checkbox"/>	2520 <input type="checkbox"/>	1800 <input type="checkbox"/>	960 <input type="checkbox"/>
Information Technology	40/-	1675 <input type="checkbox"/>	1440 <input type="checkbox"/>	1010 <input type="checkbox"/>	720 <input type="checkbox"/>	380 <input type="checkbox"/>
BenefIT	50/-	2095 <input type="checkbox"/>	----- <input type="checkbox"/>	1250 <input type="checkbox"/>	895 <input type="checkbox"/>	475 <input type="checkbox"/>
Facts For You	100/-	4190 <input type="checkbox"/>	3600 <input type="checkbox"/>	2520 <input type="checkbox"/>	1800 <input type="checkbox"/>	960 <input type="checkbox"/>

- ☒ **Don't Miss A Single Issue!**
- ☒ **Ensure Regular Supply**
- ☒ **FREE Replacement Policy***
- ☒ **Get Your Money Back if Not Satisfied****
- ☒ **Order Now For FREE Home Delivery**

* Replacement will be made upon receipt of the damaged / non-receipt copies within 30 days of its publication
** After three month, if you are not satisfied with the magazine, your balance amount will be returned.

Name _____ Company Name _____ Designation _____

Mailing Address _____ Distt. _____

City _____ Pin Code _____ State _____ Phone _____

Email _____ Subscription No. (for existing subscribers only) _____

I would like to subscribe to the above (✓)marked magazine(s), for the period _____. Please find enclosed a sum of Rs _____ by DD/MO/crossed cheque*bearing the No. _____ dt. _____ in favour of **EFY Enterprises Pvt Ltd, payable at Delhi. (*Please add Rs 50 on non-metro cheque)**

or charge my credit card ☐ VISA ☐ MasterCard Please charge Rs. _____

against my credit card No. _____ C V V No. _____ (Mandatory)

Date of Birth ____/____/____ (dd/mm/yy) Card Expiry Date ____/____ (mm/yyyy)

Signature (as on the card)



On India's 60th Independence Day in August 2007, Sun Microsystems India had a special gift for all college students in India—a unique contest, which it aptly called *Code For Freedom*.

Here is what the company had to say: “Sixty years ago, India awoke to freedom; today, we call upon the young minds of India to awaken to the clarion call of open source. The open source movement is emerging as the new world order making a sweeping impact in the technology sector across the world. Today, open source products such as Apache Web Server, Open Office and Linux have dwarfed several proprietary software corporations.

“Sun Microsystems, a front-runner in professionally backing open source and one of the largest contributors to this movement, takes this opportunity to invite the young, dynamic academic community of India to come and become a part of the exciting open source wave. We encourage you all to embrace the open source technologies with open hands and help take these communities to greater heights.

“With this vision, Sun Microsystems is happy to announce the *Code For Freedom* contest where students across India contribute to the technologies that are empowering the participation age. Participating in this contest will provide you with precious industry experience while still learning in college. And there is more. We, in turn, reward you for your valuable contribution in taking the first steps towards the open source movement.

“Come. Be a part of this phenomenon called open source.”

The contest was a roaring success with lots of high-quality contributions to open source projects. I, myself, was a participant and had the additional role of being a mentor to many prospective contributors (being the co-ordinator of the NetBeans Community Docs at that point of time).

Code for Freedom 2008

Riding high on last year's success, *Code For Freedom 2008* (CFF) [codeforfreedom.in] was launched on September 20, 2008—one of the ways Sun celebrated ‘Software Freedom Day’ this year!

Here is what the program manager, Bhuvaneshwari Panchapakesan of Sun Microsystems, India, had to say: “Sun is rising towards making the students’ life brighter

Are You Ready to Code for Freedom?

Sun calls on students to start coding.



and making opportunities available for students to participate and win through the CFF. CFF is not just a contest, but an avenue to propel the technical skills of each of the participants. This year, *Code For Freedom* is aimed at projects that would involve open source technologies. The CFF team of Sun Microsystems has been working towards making technologies and

the contest available to every corner in India.”


Contest details

The contest is open to all university undergraduate and post-graduate students from all over India. The contest format is, however, different from last year.

In 2007, the focus was on the development of Sun's open source products and participants had to contribute to open source technologies. This year, the focus is on the adoption of the products.

As part of this contest, participants will have to work on a project using Sun's open source products. Technologies that students can use include OpenSolaris, NetBeans, Java, GlassFish, MySQL, Java DB, PostgreSQL, OpenDS, Open xVM, Open Portal and some others.

Exact rules, regulations and technology details are available at the contest website at in.sun.com/communities/univ/codeforfreedom.

So, hurry up and celebrate freedom! December 31st is the last date for submitting your project proposals. **END** 

By: Amit Kumar Saha is passionate about writing and blogging. He works for Sun Microsystems, India. He blogs at <http://blogs.sun.com/amitsaha>.

Hundreds of Professional Educational Institutions Offering Professional Training Under A Single Roof

Why is EduTech The Best:

- * An event totally focused on professional education
- * A one-stop facility for students seeking admission into professional courses
- * Counselling by industry top-shots (HR and recruitments)
- * Major plan for advertising and promotion of EduTech Expo 2009
- * Powered by EFY Group's vision to boost professional education
- * Courses suitable for school pass-outs & fresh graduates

**EDUTECH
EXPO**
2009

13th - 14th JUNE 2009
PRAGATI MAIDAN, NEW DELHI



EFY Enterprises Pvt Ltd D-87/1, Okhla Indl. Area, Phase-I, New Delhi-110020; Ph:011-26810601-03; Fax: 011-26817563

INCLUDING PROFESSIONAL COURSE SUCH AS:

- Engineering • BCA & MCA • Medical & Dental • Aeronautics • Maritime • Clinical Research
- Hotel Management • Hospitality • Foreign Education • Animation & Graphics
- Computer Hardware • Computer Software • Electronics Manufacturing



The GRand Unified Bootloader Explained!

An overwritten MBR (master boot record) or boot loading errors, are things that can leave most users paralysed. The purpose of this article is to familiarise you with GRUB, the default bootloader in most modern Linux systems.

GRUB (GRand Unified Bootloader) is an advanced bootloader that is capable of booting multiple operating systems on a single machine. It can load *nix as well as other proprietary operating systems. The folks from the MS Windows platform are unfortunately ignorant about the concept of bootloaders. Proprietary operating systems like Windows often hide the background features of a system, like bootloaders, from the user.

With the help of a bootloader you can theoretically load hundreds of operating systems. Most familiar Linux distros currently ship with GRUB (Figure 1), by default. In short, GRUB is what is displayed immediately after the BIOS. It enables a user to select which OS the machine should

boot from a list, by using the arrow keys. One of the biggest benefits of GRUB is that it is dynamically configurable. Lilo is another bootloader, which was once the default and has now been depreciated by most distros.

How GRUB works

When a computer boots, the BIOS passes the control to the first-boot device—it may be the hard disk, CD-ROM, floppy disk, or Flash drive. MBR is the first sector of the hard disk and is only 512 bytes in size (Figure 2). This sector consists of code required to boot a PC. MBR consists of 446 bytes of primary bootloader code and 64 bytes of the partition table. The partition table records the information regarding the primary and extended partitions. Boot

loading is implemented in GRUB as Stage 1, Stage 2, Stage 1.5 (optional), etc. The primary bootloader area (446 bytes) contains Stage 1, which in turn directs you to Stage 2 (i.e., the *menu.lst* configuration file, which has the list of operating systems on the machine).

Installing and configuring GRUB

Most Linux distros come with GRUB by default. If your distribution comes with other boot loaders like Lilo or Syslinux, you can get the latest release of GRUB, as follows:

```
$ wget ftp://alpha.gnu.org/gnu/grub/grub-1.96.tar.gz
$ tar -xvzf grub-1.96.tar.gz -C .
$ cd grub-1.96
$ ./configure ; make
$ sudo make install
```

The next step is to configure GRUB by properly editing the *menu.lst* file. You can find the GRUB Stage 2 configuration file at */boot/grub/menu.lst*.

The next step is to add your installed operating system list to the GRUB menu. Each OS entry in GRUB will look like the following:

```
title          Ubuntu 8.04, kernel 2.6.24-17-generic
root           (hd0,4)
kernel         /boot/vmlinuz-2.6.24-17-generic root=/dev/
sda5 ro quiet splash vga=773
initrd         /boot/initrd.img-2.6.24-17-generic
quiet
```

Here, 'title' is the display name for the operating system that will appear on the GRUB bootloader screen. The following list describes what each term in the above snippet means:

- title – the display name of an operating system
- root – the partition where the kernel is located
- kernel – the path of the kernel location with specific boot parameters (space separated)
- initrd – the path of the initial ramdisk file

Going back to the snippet again, you will notice that the 'root' entry is given as (hd0,4). This is the standard GRUB naming convention, where:

- hd0 stands for the primary master hard disk
 - hd1 stands for the primary slave hard disk
 - hd2 stands for the secondary master hard disk
 - hd3 stands for the secondary slave hard disk
- Normally it would be hd0, where:
- (hd0,0) represents /dev/sda1, the first partition of the primary master hard disk
 - (hd0,4) represents /dev/sda5, the first logical partition of the extended partition inside the primary master hard disk
 - (hd1,0) represents /dev/sdb1, the first partition of the primary slave hard disk
- Other proprietary operating systems like Windows



Figure 1: The GRUB bootloader menu: select the OS you want to boot




Figure 2: The MBR

can be loaded by a process called chainloading, as follows (without specifying the kernel or other such parameters—we will only specify the partition in which it is installed):

```
title Windows Vista
rootnoverify (hd0,0)
```

BACKING UP A PARTITION TABLE FOR INSTANT FIXES

It is a good practice to make a back-up of your partition table. It will be useful to restore the partition table in the event of a corruption.

First, generate a back-up file with the partition information as follows:

```
# sfdisk -d /dev/sda > ~/partition_table.backup
```

Now, in case of a disaster, you can always restore it as follows:

```
# sfdisk /dev/sda < ~/partition_table.backup
```

Although, this is not related to GRUB in general, it's an important tip nonetheless, as this data is also stored in the MBR.



Figure 3: GRUB edit interface

```
chainloader +1
```

The following list describes what each term in the above snippet means:

- title – the display name of an operating system
- rootnoverify – the partition in GRUB notation, where the OS is installed
- chainloader +1 – enables chainloading

Now that the GRUB nomenclature is more or less covered, let us install GRUB into MBR now—run the following command as the root user:

```
# grub-install /dev/sda
```

That's it! GRUB is now your default bootloader.

The GRUB command line

The dynamic nature of GRUB helps us to alter its configuration before loading the operating system. Also, it doesn't need to reinstall the bootloader into MBR each time we make modifications to the *menu.lst* file. Working on GRUB's command line interface is similar to a bash terminal interface.

In order to switch to the GRUB's command line interface from the bootloader, press *C*. Then you will get a prompt as shown below:

```
GNU GRUB version 0.97 (640K lower / 3072K upper memory)

[ Minimal BASH-like line editing is supported. For the first
word, TAB lists possible command completions. Anywhere else
TAB lists the possible completions of a device/filename.]

grub>
```

Type *Help* to take a look at the available commands.

In order to boot a kernel, issue the following commands one by one:

```
root (hd0,0)
kernel /boot/kernel
```

```
initrd /boot/initrd.gz
boot
```

For just editing the current OS entry, select the required entry and press *E* (Figure 3). Then make the required modifications and press *B* to boot with the modified configuration.

Handling boot failures and MBR overwriting

A usual scenario all dual-boot (Linux and Windows) users face is when installing Windows after Linux; this causes MBR (the GRUB bootloader) to be overwritten by Windows. Following this, the computer straight away boots Windows, without displaying the entries for the other installed operating systems—this is why it is always advisable to install Linux after Windows. However, even if you've encountered a situation where you've lost GRUB, you can fix it easily.

Collect some GNU/Linux live CD like Knoppix and boot from it. If the live CD displays a GRUB menu, it is even easier. Press *C* to enter the GRUB command line:

```
grub> find /boot/grub/stage1
find /boot/grub/stage1
(hd0,4)
(hd0,8)
grub>
```

The output of the *find* command in the above snippet says that it has found two Linux installations on the system. Now, in order to install GRUB from either of these Linux installations, run the following set of commands:

```
grub> root (hd0,4)
root (hd0,4)
Filesystem type is reiserfs, partition type 0x83

grub> setup (hd0)
grub>setup (hd0)
Checking if "/boot/grub/stage1" exists... yes
Checking if "/boot/grub/stage2" exists... yes
Checking if "/boot/grub/reiserfs_stage1_5" exists... yes
Running "embed /boot/grub/reiserfs_stage1_5 (hd0)"... 19
sectors are embedded.
succeeded
Running "install /boot/grub/stage1 (hd0) (hd0)1+19 p
(hd0,4)/boot/grub/stage2 /boot/grub/grub.conf"... succeeded
Done.

grub>
```

That's it. Your Grub is now restored back into the MBR.

Alternately, you can boot into the live CD and get a root prompt:

```
# mkdir /mnt/fixroot
```

```
# mount /dev/sda5 /mnt/fixroot
# mount --bind /dev/ /mnt/fixroot/dev
# chroot /mnt/fixroot
# grub-install /dev/sda
```

What has to be done above is as follows: mount the root device (`/dev/sda5`) to `/mnt/fixroot`. The devices currently available to the live system are then bound to `/dev/` of the root partition (`/dev/sda5`) at `/mnt/fixroot/dev`. Finally, we temporarily change-root to the filesystem at `/dev/sda5` using the `chroot` command and execute `grub-install` to fix MBR. (Of course, don't forget to change `/dev/sda5` to the correct Linux partition on your system.)

Forgot your root password?

If you have forgotten the root password of your Linux system, there's no need to panic! The fix is quite simple. Reboot your system. At the GRUB graphical menu, press `E` to edit and add the following parameters to the kernel argument:

```
kernel /boot/vmlinuz-2.6.24-17-generic root=/dev/sda5 rw
init=/bin/bash
```

Here, by appending the `init=/bin/bash` argument to the kernel line, we are telling Linux to immediately enter a bash prompt after booting the kernel. You can now reset the root password using the `passwd` command, as follows.

```
bash # passwd
```

Now, you may wonder that if it's so simple to reset the root password, then ordinary users can use this feature to their own advantage. The next section deals with how to password protect GRUB, so that unauthorised users can't reset root passwords.

Password protecting GRUB

Generate a MD5-encrypted password for GRUB as follows:

```
[slynux@gnubox ~]$ /sbin/grub-md5-crypt
Password:
Retype password:
$1$tIwKk$K2ZwLi3kmzjssimf7K.Sh/
[slynux@gnubox ~]$
```

Now, append the MD5 hash to your `/boot/grub/menu.lst` file as follows, at the top of the file after the commented texts:

```
# menu.lst - See: grub(8), info grub, update-grub(8)
#
#      grub-install(8), grub-floppy(8),
#      grub-md5-crypt, /usr/share/doc/grub
```

CUSTOMISING A SPLASH IMAGE

GRUB usually comes with a visually appealing graphical boot menu. The background picture can be customised to your tastes, though. The Splash image is shown in the background of the GRUB bootloader screen after you switch on your PC.

In order to build a GRUB-compliant image to replace your current Splash screen, follow these steps:

1. Create an image of file format type `xpm.gz`, size 640 x 480 pixels with 14 colours only. An existing image can be converted to this format by using the `convert` command (which is a part of the ImageMagic package) as follows:

```
$ convert splash.png -resize 640x480! -colors 14 -depth 8
splash.xpm.gz
```


2. Copy `splash.xpm.gz` to the `/boot/grub/` directory. Now, edit the `/boot/grub/menu.lst` as follows by adding the following line (or replacing the text if the line already exists) before the OS specifications are listed:
`splashimage=(hd0,4)/boot/grub/splash.xpm.gz`
Reboot your computer to check the new bootloader image.

```
# and /usr/share/doc/grub-doc/.

password --md5 $1$tIwKk$K2ZwLi3kmzjssimf7K.Sh/
```

If GRUB is password protected, you won't be able to enter the edit mode by pressing `E`. Rather, you have to enter the password by pressing `P` first and `E` afterwards, to edit the menu.

What's next ?

GRUB has been undergoing mass development over the course of time—although what we just discussed is the default bootloader of most Linux distros, it has now been renamed to 'GRUB Legacy' by the GNU Project. GRUB 2, which is currently under development, is a new generation bootloader written completely from scratch, and its dynamic nature is improvised with certain GRUB moduling systems. It has various features ranging from i18n support, scripting support, cross-platform installation, etc. For more information, have a look at www.gnu.org/software/grub/grub-2.en.html. 

By: Sarath Lakshman is an 18 year old hacker and free software enthusiast from Kerala. He loves working on the awesome GNU/Linux environment and he contributes to the PiTiVi video editor project. He is also the developer of SLYNIX, a distro for newbies. He is currently studying at Model Engineering College, Cochin. He blogs at <http://www.sarathlakshman.info>

The Building Blocks Of Your Drupal Website—Part 3

The Content Construction Kit



All websites have their own ways of adding content. This requires significant effort in terms of programming and system architecture. Drupal has a scalable built-in architecture that can help you design any number of content addition forms, within minutes.

By now I'm sure you have been able to set up your Drupal website. I'm also pretty certain that you have managed to add features like blogs, books, etc, as well as configure different sections on it. In this article we will learn about the Drupal directory structure. We will also discuss user access features of a Drupal website, and then we will discover how to add our own content type to the website.

Drupal directory structure

When you install any version of Drupal, it comes with the following basic directories:

- *Themes*: This directory holds the template engine and the default Drupal themes in their own directories. More themes can be added in this directory. However, it is advised that contributed themes should be added in the *Sites* directory.
- *Modules*: This holds all Drupal core modules, where each module resides in its own directory. More modules can be added here. However, it is advised to use the *Sites* directory for new modules.
- *Scripts*: This holds shell and Perl utility scripts. It is not required in the Drupal page request cycle; however, it has scripts that are used for checking syntax, code clean-up and helps in *cron* execution.
- *Includes*: The function library used in Drupal is here.
- *Profiles*: The installation profile for a site resides here. Profile is basically a script that performs common installation and configuration tasks like enabling a few core and contributed modules that are required whenever you install a Drupal website. The Drupal Core has a 'Default' profile; however, if there are more sub-directories here with a *.profile* file in it, then when you first install your Drupal website, it asks you which profile you want to use.
- *Misc*: This directory holds JavaScript files, icons and images that are required by Drupal Core.
- *Sites*: This is a directory that holds all user files, custom, contributed themes or modules, and configuration files. By default, it has two directories: *All* and *Default*. If you want to upload a new theme, then put it in the *sites/all/theme* directory. Similarly, for a new module, use *sites/all/modules*. The default contains the Drupal configuration file *setting.php*. Figure 1 shows what the directory structure looks like.

User management

Drupal comes with basic features of user management, like sign-up, 'forgot password?', user login and user role assignment. By default, Drupal has two roles: 'anonymous users' and 'authenticated users'.

Consider a case when you want to delegate the ability to add or moderate content to other users on your site, but do not want to share all administrative abilities with them. Drupal's role management features come handy in such scenarios. Here is how you can do it.

Create a User Role: Log in with the first user of your website. Go to *Administer*→*User Management*→*Roles*. You will be taken to a page that displays a list of default user roles, viz. 'Anonymous User' and 'Authenticated User', followed by a text box and a button to add roles. Create a new role by adding 'Content Editor' in the text box and click the *Add Role* button. The page will refresh and you will see the role 'Content Editor' listed in the list with an 'Edit Permissions' link next to it in the list (Figure 3).

Assign appropriate access permissions to the user role: Stay logged in the site with first user credentials. You can assign permissions to the user roles in your website. This feature allows you to control what a user can do on your website. There are multiple ways to reach this page. First, *Administer*→*User Management*→*Permissions* allows you to see the permissions of all user roles at a glance. Second, if you navigate to *Administer*→*User Management*→*Roles* and click *Edit Permissions* next to a role, you will see the permissions of the selected user role. The *Permissions* page lists all permissions from all the enabled modules in your Drupal install.

Let's assign some permissions to the 'Content Editor' user role. If you have enabled the blog module, then under *Blog Module* click the check boxes next to the following permissions: *Create blog entries*, *Delete own blog entries*, *Edit own blog entries*. If the book module is enabled and you want your content writer to be able to post content as books, then choose the *Add content to books* permission under the book module list. Similarly, choose *Post comments without approval* under the comment module, *Access content* under the node module, *Create book content* under the node module, and finally, click *Save Permissions* (Figure 4).

You have now created a user role with appropriate permissions. Let's create multiple users to validate the same. Go to the *Administer*→*User Management*→*Users*→*Add User* page. First create the user 'guestuser' and do not assign the person a role. Again, create a user 'writer' and assign the 'Content Editor' role to this user. Now log out from the first user account and log in to the site with each user, one after another, and notice the differences in menu options available below the *My Account* link. The 'guestuser' only sees *My account* and *Log out*, while 'writer' has an additional link of *Create Content* under which the person has got permission to add blog content to the site.

Here's an exercise: Set up permission for your website's

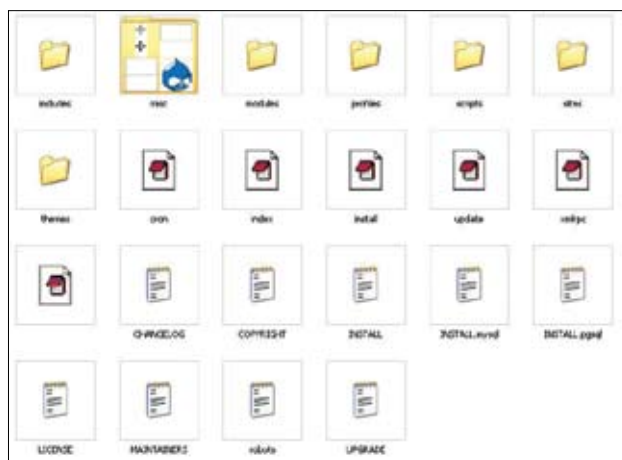


Figure 1: Drupal's directory structure

visitors to add comments on the website and access other users' profiles.

Access Rules: If you want to disallow users to sign up to the website with their e-mail address as *hotmail.com*, then you need to create an access rule. Navigate to *Administer*→*User Management*→*Access Rules*, and click *Add Rule*. Here, choose the *Access type* as *Deny*, *Rule type* as *Email*, and *Mask option* as *%hotmail.com%*, following which, click *Add Rule*. This feature helps in designing access policies that can help you control Spam on your website (Figure 3).

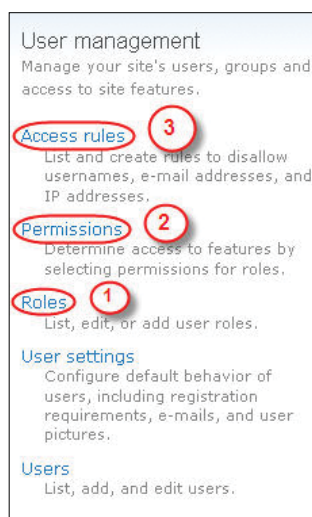


Figure 2: User management options

Designing your own content type: CCK module

The CCK or Content Creation Kit module is a powerful module and can be seen under the list of 'Must-Have' on most of the Drupal websites. This module allows a website administrator to define and configure new content types and content fields in their website. The CCK module comes with several ancillary modules that allow you to create fields of wide variety, viz., image field, e-mail field, date field, etc, apart from basic form fields. This saves you a lot in programming efforts.

The following is a list of some fields that the CCK module can produce:

- Text box and text areas
- Select and multiple select list
- Checkboxes and radio button fields
- User reference selects and auto-lookups
- Node reference selects and auto-lookups
- Date field: combos for day, month and year
- Image field: image upload and image cropping

Roles		
Name	Operations	
anonymous user	locked	edit permissions
authenticated user	locked	edit permissions
Content Editor	edit role	edit permissions
<input type="text"/> <input type="button" value="Add role"/>		
Access rules		
<input type="button" value="List"/>	<input type="button" value="Add rule"/>	<input type="button" value="Check rules"/>
Access type	Rule type	Mask
deny	e-mail	%hotmail.com%
		Operations
		edit delete

Figure 3: Roles listing and role access rules

Permissions			
Permission	anonymous user	authenticated user	Content Editor
comment module			
access comments	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
administer comments	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
post comments	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
post comments without approval	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
node module			
access content	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
administer content types	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
administer nodes	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
create book content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
create news content	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
create page content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
create story content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
delete any book content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
delete any news content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
delete any page content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
delete any story content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
delete own book content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
delete own news content	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
delete own page content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
delete own story content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
delete revisions	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
edit any book content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
edit any news content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
edit any page content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
edit any story content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
edit own book content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
edit own news content	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
edit own page content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
edit own story content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
revert revisions	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
view revisions	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
user module			
access user profiles	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
administer permissions	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
administer users	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
change own username	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input type="button" value="Save permissions"/>			

Figure 4: Permission assignment to user roles

- Media field: audio and video upload

CCK fields come with many nice features, such as AJAX auto-lookup to replace long select boxes and multiple values that enable users to enter more than one entry in a field. For example, uploading multiple images via the same image upload field.

By default, the feature to add new content type is

Modules			
<input type="button" value="List"/> <input type="button" value="Uninstall"/>			
CCK			
Enabled	Name	Version	Description
<input checked="" type="checkbox"/>	Content	6.x-2.0-rc10	Allows administrators to define new content types. Required by: Content Copy (disabled), Content Permissions (disabled), Fieldgroup (disabled), Node Reference (enabled), Number (enabled), Option Widgets (enabled), Text (enabled), User Reference (enabled)
<input type="checkbox"/>	Content Copy	6.x-2.0-rc10	Enables ability to import/export field definitions. Depends on: Content (enabled)
<input type="checkbox"/>	Content Permissions	6.x-2.0-rc10	Set field-level permissions for CCK fields. Depends on: Content (enabled)
<input type="checkbox"/>	Fieldgroup	6.x-2.0-rc10	Create display groups for CCK fields. Depends on: Content (enabled)
<input checked="" type="checkbox"/>	Node Reference	6.x-2.0-rc10	Defines a field type for referencing one node from another. Depends on: Content (enabled), Text (enabled), Option Widgets (enabled)
<input checked="" type="checkbox"/>	Number	6.x-2.0-rc10	Defines numeric field types. Depends on: Content (enabled)
<input checked="" type="checkbox"/>	Option Widgets	6.x-2.0-rc10	Defines selection, check box and radio button widgets for text and numeric fields. Depends on: Content (enabled) Required by: Node Reference (enabled), User Reference (enabled)
<input checked="" type="checkbox"/>	Text	6.x-2.0-rc10	Defines simple text field types. Depends on: Content (enabled) Required by: Node Reference (enabled), User Reference (enabled)
<input checked="" type="checkbox"/>	User Reference	6.x-2.0-rc10	Defines a field type for referencing a user from a node. Depends on: Content (enabled), Text (enabled), Option Widgets (enabled)

Figure 5: Modules available on installation of CCK

available with Drupal Install; however, to be able to add fields to your new content type you need to download the CCK module. To enable CCK modules in your website, here is what you need to do:

1. Download the latest CCK Module from drupal.org/project/cck, and unzip it in the `sites/all/modules` directory.
2. Log in as the first user, then navigate to `Administer→Site Building→Modules`. Now locate the *Content*, *Text*, *Number*, *User Reference* and *Node Reference* modules and enable them (Figure 5).

Note: Whenever you enable a module, Drupal Core always checks for its dependency on other modules, and gives you appropriate messages to enable those modules too.

Now, here's an exercise for you: visit drupal.org/project/Modules/category/88. It lists many more modules related to CCK—each can enable an exciting feature in your website. Download modules to enable date, e-mail, and URL fields capabilities in your website.

Create a new content type: For news content

Navigate to *Administer*→*Content Management*→*Content Types*; you will find a list of all content types on your website. By default, *Story* and *Page* content types are available. These are available on the *Create Content* page when you want to add content to your website. To add a new one, click on the *Add Content Type* tab, and fill the form so presented to give a name, type (should be unique) and description. Do read the descriptive text presented with every form field.

Let's now try to create a new content type with the name *News Content*, which will be used to add a section to your website. Here's how it can be done: To add a new one, click on the *Add Content Type* tab, fill the form that's presented to give a name, type (should be unique) and description.

Following are a few more settings you must look at:

1. *Submission form settings*: Drupal automatically adds two fields *Title* and *Body* whenever you add a new content type. To change the caption of these fields, click on *Submission form settings* to provide a new name to these fields. Let's change the caption for the *Title* field to 'Headline' and *Body* field to 'Detailed News'. You can also provide an explanation of this content type that will be displayed on top, whenever you go to the *Create Content* page to add content of this type.
2. *Workflow settings*: Drupal also allows you to set up workflow for the newly created content type. By default, it keeps the *Published* and *Promote to front page* options

as enabled so that content is available on the front page of the website as soon as it is published. Let's keep the default settings for now.

3. *Comment settings*: Here you can enable the content to get comments, specify who can post a comment, how the comments posted will be shown, etc. Let's keep the default settings for now; you can always edit these settings at a later stage.

Having configured these settings, click the *Save Content Type* button to save them. The 'News Content' will now be available on the content type list page.

Add a field to 'News Content': On the content type list page, the *News Content* type will be displayed with three options, viz., edit, manage fields and delete. Click on *Manage fields*. You will be taken to a page that lists all the fields available in this content type with an option below to add a new field to the content type. Let's add a new field for the reporter. We can give a label 'Reported by' and call it 'reporter' of type text and save it. You will be taken to another page asking for more information on the field. The information relates to field size, its default value, its maximum length, etc.

Next, you will be taken to the *Manage fields* tab of the newly added content type, 'News Content' that lists all its fields. You can change the order of the fields as you want them to appear on your page. Let's move our 'Reported By' field next to 'Headline' and save the settings. You will notice that the *configure* and *remove* options are available against

WE TOPPED AGAIN



BEST CERTIFIED

RED HAT TRAINING PARTNER

SOUTH REGION (2006-07)



BEST CERTIFIED

RED HAT TRAINING PARTNER

SOUTH REGION (2005-06)



RED HAT LINUX

**RHCE / RHCT / RHCSA TRACK
RHCE EXAM ORIENTATION**



Complete Open Source Solutions

#512, Aditya Trade Center, Ameerpet, Hyd- 38
Tel: +91-40-66773365, Fax: +91-40-66463365
web: <http://www.cossindia.net>

Simply Believe

NO Server Setup Required

Easy connection to Linux, Windows, Citrix, Goglobal etc. Servers

Embedded Openoffice Suite, Email Client (Thunderbird), Web browser (Firefox), Media Player (Audio/Video), PDF Reader, Messengers, Skype

Embedded DOS Client (connect to any dos folder on Linux, Novel, Windows, Server) - (Full Screen Dos, Switchable session)

Multi-Session, Multi-Protocol Multi-server support

Local Web Browser (with JVM 1.6, flash plugins, media codecs)

Enterprise Class ThinClient OS



Highly Efficient ThinClient Embedded OS

With a Complete range of ThinClient Pcs

An ISO Certified Company
Enjay Network Solutions

www.enjayworld.com | info@enjayworld.com | support@enjayworld.com
Phone : 0260-3251732, 3241732, 3203400, 2785124 | 09898007650, 09377107650

Content types List [Add content type](#)

To create a new content type, enter the human-readable name, the machine-readable name, and all other relevant fields that are on this page. Once created, users of your site will be able to create posts that are instances of this content type.

Identification

Name:
The human-readable name of this content type. This text will be displayed as part of the list on the create content page. It is recommended that this name begin with a capital letter and contain only letters, numbers, and spaces. This name must be unique.

Type:
The machine-readable name of this content type. This text will be used for constructing the URL of the create content page for this content type. This name must contain only lowercase letters, numbers, and underscores. Underscores will be converted into hyphens when constructing the URL of the create content page. This name must be unique.

Description:
A brief description of this content type. This text will be displayed as part of the list on the create content page.

[Submission form settings](#)

[Workflow settings](#)

[Comment settings](#)

[Save content type](#)

Figure 6: Add Content Type screen

the 'Reported By' field, but the same is not available for the 'Headline' and 'Detailed News' as these two are the default and mandatory fields required for every content type (Figure 6 and 7).

Note: Once you add a field to any of the content types, the same would be available for inclusion as the field in any other content type also. All available fields can be found on the *Administer*→*Content Management*→*Fields* page. Hence, a field once added can be use by all available content types if they require.

Adding content using 'News Content' content type:

Since you are logged in as the first user, you will find the 'News Content' option under the *Create Content* menu option. Try adding some news using this content type. At this point if you log out and log back in again with either 'guestuser' or 'writer', you will not find this content under *Create Content* option in their log in. You need to give permission to the user role of these users if you want them to be able to add the 'News Content' in your website. Following section describes how:

Permission for new content type: Whenever you create a new content type, a few permissions are automatically created under *Administer*→*User Management*→*Permissions*. These can be found under the *Node* module section in the

News Content Edit [Manage fields](#) [Display fields](#)

[Saved field Reported By.](#)

Add fields and groups to the content type, and arrange them on content display and input forms. Note: Installing the Advanced help module will let you access more and better help.

Label	Name	Type	Operations
Headline	Headline module form.		
Detailed News	Detailed News module form.		
Menu settings	Menu module form.		
Reported By	field_reporter	Text	Configure fields

Add

[New field](#)

Label
Field name (a-z, 0-9, _)

[Select a field type](#) [Select a widget](#)

Type of data to store. Form element to add the data.

[Save](#)

Figure 7: The Content Field Addition screen

page. These permissions are to *Create*, *Delete any*, *Delete own*, *Edit any* or *Edit own* content for the 'News Content' content type (Figure 4). If you want someone other than the first user to add news content to your website, you need to assign these permissions to a user role and also assign this role to the user whom you want to delegate the permission to.

Log in to the website as 'writer'. You will notice that the 'News Content' option is not available for this user. Next, assign the permission to *Create*, *Delete own*, and *Edit own news content* to the 'Content Editor' role from the first user login. Again, log in with 'writer' and notice that now the user can add content on the website.

writer

- [My account](#)
- [Create content](#)
- [Blog entry](#)
- [News Content](#)
- [Log out](#)

Writer with Content Editor Role

guestuser

- [My account](#)
- [Log out](#)

user with no permission to add content

Figure 8: User with appropriate permissions can add content

Wrapping up

In this instalment, we discussed the user management features of Drupal that can help you delegate limited abilities to a user role on your website. We also discussed how to add new content types to your website. Tune in next month to learn some more advanced and nifty features in Drupal Theming and the ways to present data from your new content types. **END**

REFERENCES:

- *Drupal CCK Modules page:* drupal.org/project/Modules/category/88
- *Learn more about CCK:* drupal.org/handbook/modules/cck

By: Dheeraj Dagliya is director, Gloscon Solutions Pvt Ltd, Ahmedabad, a 100 per cent EOU. Gloscon is an open source product development company specialising in Drupal and Ruby on Rails (RoR) frameworks. The author can be reached at drupaldesk@gloscon.com

Festival Season

Celebration Begins with Offers from

EFYbonanza
Tech Products. Amazing Offers. •COM

EFY Group's latest initiative to bring to you innovative, hi-tech products, and that too with special discounts

To Buy, Call: 9312888878

Or visit: www.efybonanza.com
To Buy or Get More Details
+91-11-26810601/02/03
To Sell Your Products through EFY Bonanza Call



SVB Digital Photo Frame Keychain
Rs. ~~1,500/-~~ Rs. **1,100/-**



Wow Wee SpiderSapien Robot
Rs. ~~4,000/-~~ Rs. **3,700/-**



Portable Media Player with in-built Camera
Rs. ~~4,400/-~~ Rs. **3,300/-**



Wireless Camera GZ 608
Rs. ~~2,750/-~~ Rs. **2,500/-**



Magicon Mobile Phone ENVY MG407
Rs. ~~5,599/-~~ Rs. **5,000/-**



HitPlay Clocky Alarm Clock : Chrome
Rs. ~~3,299/-~~ Rs. **3,000/-**



Button Type Cordless Spy Camera
Rs. ~~2,199/-~~ Rs. **2,000/-**



Magic Bullet 17 Piece Food Processor Blender Juicer
Rs. ~~2,999/-~~ Rs. **1,900/-**



Hitplay Personal Weight Tracking Scale
Rs. ~~2,599/-~~ Rs. **2,400/-**



SVB Digital Photo Frame Round Shape
Rs. ~~2,000/-~~ Rs. **1,400/-**



All-in-one Mobile Phone Charger
Rs. ~~499/-~~ Rs. **400/-**



SVB USB Speakers 289
Rs. ~~540/-~~ Rs. **400/-**



Axiom Car Inverter 150 Watts
Rs. ~~1,990/-~~ Rs. **1,800/-**



Maps of India CD version 6.0
Rs. ~~995/-~~ Rs. **900/-**



DVD Player with 7 inches LCD
Rs. ~~6,990/-~~ Rs. **5,000/-**



Car Washer
Rs. ~~5,799/-~~ Rs. **4,500/-**



800 Watts Vacuum Cleaner
Rs. ~~1,290/-~~ Rs. **1,000/-**



Apex Magic Power Saver Switch (PIR Sensor) - 360 Degree
Rs. ~~1,700/-~~ Rs. **1,000/-**



Mitashi KNM-2GB MP4 Player
Rs. ~~4,990/-~~ Rs. **4,700/-**



Car Vacuum Cleaner With Powerful Search Light
Rs. ~~999/-~~ Rs. **600/-**



Dome Camera With Audio
Rs. ~~1,960/-~~ Rs. **1,700/-**



SVB USB Steering Wheel
Rs. ~~2,400/-~~ Rs. **1,800/-**



SVB Laptop Car Charger
Rs. ~~1,400/-~~ Rs. **750/-**



Mitashi Home Theatre PH-6050
Rs. ~~3,590/-~~ Rs. **3,200/-**



Electronic Sewing Machine
Rs. ~~2,999/-~~ Rs. **1,600/-**



MX Universal Travel Adapter - MX-2725A
Rs. ~~904/-~~ Rs. **700/-**



Bens Induction Cooker
Rs. ~~2,490/-~~ Rs. **1,800/-**



Hitplay Sound Activated T-Shirt
Rs. ~~1,999/-~~ Rs. **1,300/-**



Citizen Digital Automatic Blood Pressure Monitor
Rs. ~~1,999/-~~ Rs. **1,900/-**



Hitplay Kitchen Scale
Rs. ~~1,599/-~~ Rs. **1,500/-**



SVB Mini USB Paper Shredder with letter opener
Rs. ~~980/-~~ Rs. **750/-**



Blood Glucose Monitoring Kit : GlucoLeader Value
Rs. ~~1,995/-~~ Rs. **1,350/-**



Body Fat and Water Weighing Scale
Rs. ~~2,500/-~~ Rs. **2,100/-**



Apex Video Door Phone with Night Vision
Rs. ~~7,999/-~~ Rs. **6,000/-**



eLanguage Educational Software Kit : French
Rs. ~~999/-~~ Rs. **900/-**

The EFY Assurance: • We search for innovative hi-tech products • We get special pricing for our readers • We tie-up with their importers or manufacturers • We facilitate the deal between you and them • We keep your money secure and release it to them only after confirming receipt • We refund money if the product does not reach you within a few working days

An **EFY** Group Website

| EFY Enterprises Pvt Ltd, D-87/1, Okhla Industrial Area, Phase-1, New Delhi-110020

Linux Threads



How to use threads while programming.

Maximum utilisation of the CPU has always been the main goal of computer software and systems designers. Threading is one software technique that not only achieved this but also made software systems more and more modular. After conforming to the POSIX standards, thread-based architectures have proved to be easily portable and prolific. Linux supports threads both at the user and kernel level. This article focuses on Linux threads at the user level.

Basic concepts

From an operating system perspective, multiple 'execution contexts' or 'tasks' need to be created to maximise CPU utilisation. So, if you wait for some resource or I/O, the scheduler schedules another task so that the CPU does not waste any of the cycles. Creating such a 'task' or 'execution context' (or 'process,' in Linux terminology) is generally an expensive operation, because it involves the replication of the complete address space as well as replicating all the resources (such as open file descriptors, etc) by the parent 'process'. One more setback also occurs at the system level when one or more 'processes' want to talk. They can talk only through some kind of FIFO/pipe/socket or 'mmap', and all of them either involve the copying of data from the user space to kernel space and vice-versa, or the use of page table manipulation.

The operating system designers have thought of one more way to solve this problem with a lower penalty,

and that's where the concept of threads comes in. The creation of threads also means creating separate 'execution contexts', but with less penalty.

Threads share common 'data sections', 'heap' and system resources such as the environment, file descriptors, etc, and hence the creation and deletion of a thread is very 'lightweight' as compared to the creation of processes.

Since the 'process address space' is shared between the threads in a given process, the communication between threads also does not have as many penalties.

Figure 1 depicts this point more clearly

Thread models

Once you have decided that the implementation will be done using threads, the first task is to logically divide the multiple functionalities in a given system so that a separate thread could be assigned to perform each individual function.

The Master-Slaves model: Consider an ATM machine system. Here, one thread could be made to manage the touchpad/keypad, another could manage the display/beep functionalities, a third could manage the remote connection with the bank server along with transactions and actual money vending. Now a central manager/master thread will be present to keep all of them synchronised. This is precisely what the 'Master-Slaves model' suggests.

The pipeline model: Here the concept is that the first thread runs and produces some outputs based on the

system inputs. These intermediate outputs are handed over to the second thread, and so on. The last thread delivers its output as the system output. Here the previous thread becomes free for the next input as it hands over the output to the next thread in the pipeline. Typically, this model is used in computation intensive applications.

The background task model: Consider the case of an operating system at start-up. It generally needs to initiate a DHCP query to acquire the IP address from the network as a part of its boot process. Now this start-up operation cannot wait till the system acquires the IP address and hence could choose to spawn a thread for this activity and proceed ahead with the next start-up sequence.

These models are generally used across the industry. The choice of the right model is mostly specific to a given problem and needs proper analysis. The important consideration here is to try and reduce the communication and switching between two threads as much as possible, for better system performance.

Manipulation on threads

After we chose the right 'model', what remains is implementation.

The following are two important structures that you need to manipulate when playing around with threads in Linux:

- `pthread_t`: Generally used as a handle to a thread.
- `pthread_attr_t`: Contains thread attributes, i.e., the parameters that control the behaviour of a given thread.

It is not recommended that you populate them directly, but a number of APIs are provided by the 'pthread' library (or NPTL library) to manipulate them.

The following listing states a few of them:

- `pthread_attr_init()` and `pthread_attr_destroy()`: Allocates and destroys the structure.
- `pthread_attr_setdetachstate()` and `pthread_attr_getdetachstate()`: Sets/gets the 'detachstate' attribute (whether the thread can be joined on termination or terminated independently).
- `pthread_attr_setschedpolicy()` and `pthread_attr_getschedpolicy()`: Sets/gets the 'schedpolicy' attribute (FIFO scheduling/round-robin scheduling).

Creation and deletion: `pthread_create()` and `pthread_exit()` are the two basic APIs that could be used for this purpose. They are fairly simple to use and hardly need any explanation. Other APIs that could be used are:

- `pthread_join()`: It blocks the current thread until another one terminates. It partially achieves synchronisation.
- `pthread_cancel()`: It cancels the execution of another thread.

Synchronisation: Being in a multi-threaded environment, there's no guarantee when a given thread would be invoked unless a proper 'synchronisation' mechanism is deployed. Synchronisation basically has two aspects. One is to make one thread wait until the other

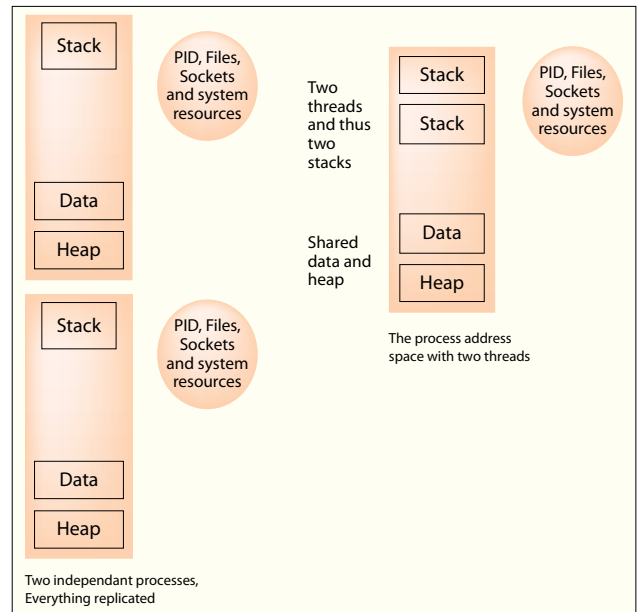


Figure 1: Process versus threads

has completed a specific activity. Here I intend to highlight the scheduling aspect of the threads. The second aspect of synchronisation is related to resource sharing between multiple threads. Here, for the caller it is not so important which thread grabs the resource first, but that the resource should be in a consistent state before it is taken by another thread. Yes, I am talking about 'mutual exclusion' here. Going back to implementation, the following are two different types of interfaces that are provided by the 'pthread' library (or NPTL library):

1. *Synchronisation using condition/signal (for execution synchronisation)*
 - `pthread_cond_init()`: Initialises a 'condition' variable.
 - `pthread_cond_destroy()`: Destroys a 'condition' variable.
 - `pthread_cond_signal()`: Signals a 'condition'.
 - `pthread_cond_wait()`: Waits on a 'condition'.

The use of these APIs is fairly simple and can be easily understood by reading their manuals.
2. *Synchronisation using mutex (for protection synchronisation)*
 - `pthread_mutex_init()`: Initialises the 'mutex lock' variable.
 - `pthread_mutex_destroy()`: Destroys a 'mutex lock' variable.
 - `pthread_mutex_lock()`: Attempts and acquires a 'mutex lock' (blocking call).
 - `pthread_mutex_trylock()`: Attempts and acquires a 'mutex lock' (non-blocking call, one attempt).
 - `pthread_mutex_unlock()`: Releases the 'mutex lock'.

These APIs too are fairly simple and only require the man pages, to get you familiarised.

Thread local storage: It enables two threads referring to the same static or global variable that refer to different memory locations, thereby making the variable 'thread local'.

A key or an identifier needs to be created for such variables. It could be done using `pthread_key_create()`. `pthread_getspecific()` and `pthread_setspecific()` are used to assign/ fetch the 'thread specific' address assigned to 'key'. The corresponding clean-up function is `pthread_key_delete()`.

GCC extension: Thread local storage can also be achieved in GCC by the virtue of additional storage class keyword '`__thread`'. For example:

```
__thread int ThreadLocVar;
```

Some pitfalls

- Threads in abundance can give rise to thread switching and thread management penalties. This is something that needs to be watched for at the design time.
 - The logic for mutex/conditions needs to be designed carefully. A small mistake there can cause mutex/conditions deadlocks or starvations.
 - Writing 'Thread Safe' code is extremely important. It basically boils down to understanding what is shared and what's not shared. Following are the things that are shared across threads (and hence need managed access):
 - Global data
 - Static variables.
 - Heap/dynamically allocated memory if the address gets communicated across threads.
 - Any system resource (such as file descriptors, sockets etc.)
- Things that are not shared:
- Local and stack variables.
 - Thread specific data.
- Consistent system state and proper recovery in case of thread getting killed.

What happens in Linux kernel?

When `pthread_create()` gets invoked from user space, the ultimate function that gets invoked in kernel is '`clone()`', which is the same function that is used to create a copy of the calling process. (It also gets called as a result of `fork()`, but with a different set of arguments.)

These 'execution contexts' created by '`clone()`' are schedulable entities for the Linux kernel. Hence both 'LinuxThreads' and 'NPTL' are 1:1 implementations, which means that each thread maps to a kernel scheduling entity.

What went wrong in LinuxThreads?

Multi-threading support has been brought to Linux by the LinuxThreads project. Although a pioneer, the LinuxThreads implementation had some disadvantages, particularly in the areas of signal handling, scheduling, inter-process synchronisation, performance on SMP architectures, and the fact that it did not conform to POSIX.1 specifications. Some of the shortcomings with LinuxThreads were:

- LinuxThreads had a dependency on the 'manager

thread' to manage (allocate/de-allocate) and perform signal related functions.

- 'Manager threads' incur the additional overhead of thread switching.
- 'Manager threads' could run on only one CPU and hence had scalability issues on SMP and other multi-processor architectures.
- Each thread had a separate process ID, which was a non conformance to POSIX specs.
- Because each thread is a process, there was a (system dependant) limit on the number of threads that could be created for an application.
- Signals are delivered to individual threads rather than to a process. For example, if SIGSTOP is sent by GDB, it just stopped the relevant thread rather than the complete processes.

NPTL: The way to go!

NPTL is the acronym for Native POSIX Threads Library. This is the modern *Pthreads* implementation. Compared with LinuxThreads, NPTL provides closer conformance to the requirements of the POSIX.1 specifications and better performance when creating large numbers of threads. The following are salient features of the NPTL:

- There is no concept of 'manager threads'. The important functionalities of the manager threads have been pushed into the kernel. This makes the NPTL more scalable and architecture independent.
- Each thread returns the same process ID and hence the signal given to a process (like SIGSTOP by GDB) reaches all the threads in the process.
- NPTL introduces a new mechanism called a *futex* (Fast Userspace muTex). It provides yet another means of synchronisation with hardly any context switch.

So far we have covered a brief introduction to threads, in general, along with some design and implementation perspectives. We also discussed a bit about the past, present and future of 'threads' on Linux. Of course, this is just a start. Now, let's start thinking 'parallel'...  **END**

REFERENCES:

- Manual page: linux.die.net/man/7/pthreads
- Thread FAQs: pauillac.inria.fr/~xleroy/linuxthreads/faq.html
- Thread usage: www.yolinux.com/TUTORIALS/LinuxTutorialPosixThreads.html#PITFALLS
- Interview with thread guru David Butenhof: www.thinkingparallel.com/2007/04/11/ten-questions-with-david-butenhof-about-parallel-programming-and-posix-threads

By: Nilesh Govande. The author is a Linux enthusiast and could be contacted at nileshgovande@yahoo.com. His areas of interest include Linux system software, application development and virtualisation. He is currently working with LSI Research & Development Centre, Pune.

BraveGNUIndia

Building Skills, the Delhi Way



FREDERICK NORONHA

Encouraging localisation to take Free Software to the masses.

Debyan Banerjee of NIT Durgapur, dreams of building an active Indic OCR community. He suggests that we should not be satisfied with software that can recognise Indic characters with 95 per cent accuracy. So, why not a Web-based interface for the application, where people can upload their scanned documents and get OCR'd text instantly?

Amit Kumar of Panipat has an idea to create a GNU/Linux Live CD kiosk. Such a tool, he says, could be deployed in rural projects, for education, at shopping malls, museums, and to display campus maps. He proposes to turn PCs into a kiosk of sorts, built on a platform based on the GNU/Linux Live CD.

From elsewhere across India, the Project Hindawi team came up with ideas on how to take their concept further. Hindawi describes itself as a "...free, open-source, productised and released, complete non-English-based systems programming platform supporting all paradigms of programming languages—from assembly language to logic and functional programming."

Its laudable goal is to shatter the language barrier, allowing non-English literates to take up computer sciences at all levels, from primary school education to robotics and supercomputing, in their mother tongue, without the need to master English.

These are just three of the projects that came up at the PRC-Sarai network recently.

What's *that*? Sarai is a Delhi-based non-profit organisation that bought into the idea of creating a small-grants project. They call it the FLOSS Fellowships (and thanks for not subsuming Free Software in Open Source!).

Sarai is linked to the Centre for the Study of Developing Society, located at Rajpur Road in the Civil Lines area in North Delhi, close to North Campus, Delhi University.

While other non-profits and NGOs use donor money to create pilot projects—that are non-replicable at best or a waste, at worst—Sarai has opted for an interesting model. What this organisation does is to encourage mostly youngsters to come up with that bright idea. Selected ideas are given funding, and then, being based on Free/Libre and Open Source Software, would naturally go into a shareable software pool.

In one stroke, a number of targets are met. First, students get to build their skills at a crucial stage towards building their career path. Second, money is well spent, because it goes into building scalable skills. Third, the end

result is something that could be used again, and again, and again... because of its sharable nature.

This has been going on for some time now. This year, funding came from the National Internet Exchange of India (NIXI), www.nixi.in. The goal was to "...promote Indian language computing, and the network technology needs for India in the FLOSS domain." ILUG-Delhi was also the technical partner.

Over a six-month period, FLOSS Fellows would get Rs 70,000 or so to work on their selected projects. Interestingly, in keeping with the goal of openness, applications for this process were also put out via a public mailing list, called PRC (Project Resource Centre).

Instead of applicants gaining from the 'element of surprise', by sharing their applications they are probably also sharing inspiration across a wider area.

What is also interesting is that this attempt has already got in some useful results from the previous


years. There's the Hindi automated speech recognition portal [sourceforge.net/projects/hindiasr], NewsRack [newsrack.in], and Hindawi [www.indicybers.com].

Personally, as a journalist, I've

encountered and made use of NewsRack—a kind of Google Alerts for India, only far more relevant. One is very thrilled by what the young man behind it, Subbu Shastry, achieved. Of course, the funding and recognition probably spurred on his earlier intentions.

Sarai deserves praise. And I'm not saying this just because I was one of their early fellowship-awardees, at a time when the ideas of tapping FLOSS were not yet on the horizon.

For a change, this small organisation has shown how some funding could go a long way in taking ideas ahead, building skills and solutions, and then putting the same back into the community. It's time we worked on more innovative ways of promoting Free Software in India, instead of merely lamenting our lack of deeper involvement in the field.

See: www.sarai.net 

Frederick Noronha (FN) is an independent journalist based in Goa. See his blog at fredericknoronha.wordpress.com and encounter his links on Twitter at twitter.com/fn

GUEST COLUMN

Programming in Python for Friends and Relations



Python Scripts: Generators Make It Easier to Rescue Systems

Automate system recovery by writing small scripts.

It wasn't quite the blue screen of death. The system wanted me to fix some errors on the disk partition manually. I was uncomfortable. Just two weeks earlier, the same thing had happened. The *fsck -y* command had fixed all the errors! At the end, the only visible directory was *lost+found*.

Even if I booted from another partition, I couldn't mount the partition and take a back-up till I cleaned up the partition with *fsck*. I had already changed the SATA data cable on the previous occasion, which seemed to have fixed the problem. At the suggestion of the hardware supplier, I changed the power cable as well. The system seemed more stable now. It was no longer grinding to a

near halt with a log reporting: “ata4: hard resetting link”. However, it was too soon to rejoice.

I finally gave in and ran *fsck*. It cleaned up quite a few files/directories. It booted with errors and X wouldn't run. It was just the system partition with nothing more than the Fedora 9 installation. I had added a fair number of additional packages. It probably would have been faster to just reformat the partition and reinstall the OS. This time, I had taken the precaution of caching the downloaded RPMs on a different partition! So, the 24-hour download time would not be needed for the updates and the additional packages.

However, it seemed that this was an interesting problem. Could I recover a system which was so badly trashed? Based on the problems noticed, I used *rpm -V* on some packages and found that some libraries were missing. Some packages were trying to access information beyond the partition. To make matters worse, I had been in the middle of an update (to which my wife would say—“When are you *not*?”).

The first step was to at least measure the scope of the problem. I took a list of all the packages installed:

```
rpm -qa > installed.list
```

I wasn't about to manually verify each one of the 1,500 or so packages! So, a small Python script would be useful:

```
import os
f = open('installed.list')
fbad = open('bad_rpm.list', 'w')
for line in f.readlines():
    if os.system("rpm -V " + line[:-1]):
        fbad.write(line)
```

It was a relief to know that only about 400 packages were in a damaged state! Even this was too large a number to handle manually. Surprisingly, there were some version issues. This turned out to be because there were multiple entries for some packages, thanks to the failed update.

Write scripts—you may need them again

Fortunately, I had written a utility over a year ago to solve that problem. A combination of power failure and the UPS system battery running down in the middle of an upgrade had left my system in an inconsistent state.

The utility pieces were as follows. I'd created a dictionary of package names with attributes like the version, release, arch, etc, which would help identify duplicates.

```
import rpm
def get_packages(ts):
    mi=ts.dbMatch()
    packages = {}
    for hdr in mi:
        name = hdr['name']
```

```
        attr=[hdr['version'],hdr['release'],hdr['arch'],hdr['distribution']]
        if name in packages:
            packages[name].append(attr)
        else:
            packages[name]= [attr]
    return packages
ts = rpm.TransactionSet()
packages = get_packages(ts)
```

Especially on an x86-64 architecture, a package with the same name may occur for i386 architecture as well. Hence, that is not a duplicate. You need to check whether there are any duplicates for each name and arch combination.

```
ARCHS=['x86_64','i386','i686','noarch']
def get_duplicates(packages):
    duplicates = {}
    for name in packages:
        for arch in ARCHS:
            dups = chk_dups(packages[name],arch)
            if dups:
                duplicates[(name,arch)] = dups
    return duplicates
duplicates = get_duplicates(packages)
```

The actual work of checking duplicates is done in *chk_dups*. Assume that there is only one package with the maximum version.

```
def chk_dups(pkgs,arch):
    dup_pkgs = filter(lambda x: x[2] == arch, pkgs)
    if len(dup_pkgs) > 1:
        max_version = max([(x[0],x[1]) for x in dup_pkgs])
        newPkg = filter(lambda x: (x[0],x[1]) == max_version,
dup_pkgs)
        restPkg = filter(lambda x: (x[0],x[1])!= max_
version,dup_pkgs)
        return newPkg,restPkg
    else:
        return None
```

If I were writing this program today, I would have avoided the filter function and used list comprehension instead. For example:

```
newPkg = [ x for x in dup_pkgs if (x[0],x[1]) == max_version]
```

While I could have deleted the RPMs in the program, I felt more comfortable getting a list of duplicate package names and then deleting them from the command line.

```
def delete_duplicates(dups):
    f=open('deleteList.txt','w')
    for name in dups:
        for rpm in dups[name][1]:
```

```
rpmname = name[0] + '-' + rpm[0] + '-' + rpm[1] +
\.' + rpm[2]
f.write(rpmname + '\n')
f.close()
delete_duplicates(duplicates)
```

Now, as the root user, I ran:

```
rpm -e `cat deleteList.txt`
```

Having deleted some packages, I needed to get a fresh list of the installed packages and those that failed the verification.

The next step was to reinstall all the packages with problems. Since the RPMs were in various subdirectories of */var/cache/yum/*, I collected all of them in */opt/yum/RPMS/*. The script used was:

```
import os
LOC = '/opt/yum/RPMS/'
packages = os.listdir(LOC)
f = open('bad_rpm.list')
for line in f.readlines():
    fn = line[:-1] + '.rpm'
    if fn in packages:
        os.system('rpm -Uv -force ' + LOC + fn)
    else:
        print fn, " Not Found"
```

Some downloaded packages were lost. So, the final step was to use *yum update* to update the missing packages.

On the first occasion I had to reinstall from scratch, it had taken me well over two days to fully recover. Most of the time was spent downloading updates and packages not on the distribution DVD. Partly, it is hard to remember all the additional packages installed. My memory was often triggered by a high-interrupt from my wife—for example, “Where’s Sylpheed?”

This time, I recovered the system in little over a day; with more than half the time spent in figuring out the issues and developing the code. But now if the system winds up in the same state, I am sure I can recover in much less than half a day.

Actually, I will recover much faster because I now have a dual-boot system. I bought another disk and have a fully configured installation on that disk as well.

Postscript—a solution using generators

I recently came across an excellent presentation on using generators at www.dabeaz.com/generators. I realised that I had created temporary intermediary lists or dictionaries in order to ensure the code was easier to follow. How would the programming for fixing the issue of duplicate RPMs be different if I approached it from the perspective of generators?

I want to iterate over each package that is a duplicate and then take action on it. So, let us just create a list of them. The code needed is:

```
delete_list = []
for package in duplicate_packages():
    delete_list.append(package)
print delete_list
```

The function *duplicate_packages* looks, feels and behaves like an iterator.

If you iterate over each package, you can determine which package is a duplicate. Let us examine the header of each package. A package will be identified by the name and arch pair. The unique version is determined by the version and release pair.

```
def duplicate_packages():
    packages = {}
    for hdr in package_headers():
        key = (hdr['name'], hdr['arch'])
        version = (hdr['version'],hdr['release'])
        if key in packages:
            yield get_older(packages, key, version)
        else:
            packages[key]= version
```


The keyword *yield* has converted this function into a generator; so, we can iterate over duplicate packages. The method *get_older* is straightforward:

```
def get_older(packages,key, version):
    prev_version = packages[key]
    if version > prev_version:
        packages[key] = version
    version = prev_version
    return (key, version)
```

The method *package_headers* is another generator:

```
import rpm
def package_headers():
    ts = rpm.TransactionSet()
    mi=ts.dbMatch()
    for hdr in mi:
        yield hdr
```

The fascinating thing is that this code looks flat even though it is equivalent to nested code. It looks cleaner and is shorter.

Unfortunately, I have to wait for the system to have problems before I can test it properly. Or as some weird laws of nature go—now that I have the back-ups, I may never get a chance to use them! 

By: Anil Seth, consultant, seth.anil@gmail.com

FreedomYug

My TV Station



NIYAM BHUSHAN

Would you like to launch your own live TV station?

If you're between the ages of 21 and 35, you should be quite excited to discover you can finally launch your own live video to a global audience. If you're above the age of 35 and a senior honcho in traditional TV and broadcasting, you should be either worried, or excited; or both. A small handful of geeks have just launched *GISS.tv*.

All you need for your personal TV station is a commodity video camera, or even a Web camera attached to your networked laptop, and you can start broadcasting immediately. All this for free, and more importantly, with freedom. But first, let's clarify one thing: *GISS.tv* is not YouTube, which only allows you to publish pre-recorded and short video-segments.

Blissed-out TV

With *GISS.tv*, you can show live events as they happen, and you can go on-air with your own live news broadcasts, as well as hold live talk-shows and interviews. The possibilities are endless. How much will this cost you? Nothing! It is free as in free-of-cost. Just pay for your regular broadband connection. *GISS.tv* provides you with a free and instant sign-on channel so you can broadcast right away. Viewers are presented with a constantly-updated world-map, where they may hover their mouse-pointers over markers to discover audio and video broadcasts that might interest them.

True to the spirit of FOSS, *GISS.tv* diligently avoids all proprietary software, especially Flash, and allows viewers to watch live video using nothing more than any ordinary Java-enabled browser. Authors are encouraged to use 100 per cent FOSS file-formats such as Ogg Theora to publish their videos.

Yves and Know


Hardly ten of us showed up at Sarai in Delhi, to watch Yves Degoyon, one of the main project-leaders of *GISS.tv* demo his astounding software. Ironical for a country that boasts of more than 60 million TV-viewers. It took all of 15 minutes for Yves to connect and broadcast live video globally from his laptop and armchair. Unlike traditional TV, we could track users as they tuned in, noting their locations, and even their browser and software details. Users can also subscribe to and track your channels. Handling one video-

stream was okay, but could *GISS.tv* handle a traditional TV production environment? By using a simple drag-and-drop, Yves had scrolling text and ticker-tapes across the bottom, or any part of the screen, scrolling in any direction and at any angle. With a few more drag-and-drops, he demo-ed the handling of multiple video streams, both live and pre-recorded. Just like TV. He also got picture-in-picture, and deftly managed dissolves and transitions. *GISS.tv* can take on pretty much anything a traditional TV control-room can handle. You can even test-launch your TV station, without installing software on your machine. All you need is the Live CD, that boots your regular Windows-based PC into GNU/Linux, runs all the software, and shuts down, without installing anything on your computer.

Geeky bling-bling

The problem with *GISS.tv* is that its website and software look intimidatingly geeky. It totally misses out on the glamour and chic of YouTube or of *blip*.

tv. The excitement lies in the paradigm that will disrupt media, new media and business-models. Apart from insisting that all content must be either your own copyright or else be copyleft, with no commercial advertising, the site also states: "No racism, xenophobia,

sexism, homophobia, religious hatred or missionaryism...". However, Yves chuckles as he admits you could broadcast FOSS propaganda if you wish. If you do not want any censorship, you are welcome to use and modify the free software, set up your own servers, and go live on your own. This is where *GISS.tv* truly scores. It is no longer just about free software. It is about free media. In the 21st century, firing a shot at neurotic mass media, and its unholy nexus with politicians and business, will be one of the lasting impacts of free software. **END** 

"The excitement lies in the paradigm that will disrupt media, new media and business-models."

GUEST COLUMN

Inspired by the vision of Osho. Copyright September 2008: Niyam Bhushan. freedomyugs at gmail dotcom. First published in LinuxForYou magazine. Verbatim copying, publishing and distribution of this article is encouraged in any language and medium, so long as this copyright notice is preserved. In Hindi, 'muft' means 'free-of-cost', and 'mukt' means 'with freedom.'

MComplex Mathematical

How to use Python for typical mathematical operations.

Mathematics is the foundation of many complex computing applications that are in use today. Computer networks, information retrieval systems, search engines, game development, graphics, image processing, grid computing, security consulting, application and system software are some areas where mathematics is widely used. Meanwhile, when it comes to the Python programming language, Google, NASA and many other scientific organisations are using it significantly for various purposes as Python supports complex mathematical operations efficiently.

Python supports algebraic functions, notations, mathematical functions, simplifying and solving functionalities, etc, to solve complex equations in scientific computing. Math library, NumPy and SciPy packages can be used for computing. Math library supports mathematical functions defined by the C standard. These return float values except for some rare circumstances. NumPy is used for multi-dimensional array storage. SciPy is used for scientific programming in Python. You can download the latest versions of NumPy and SciPy from www.scipy.org.

The math and cmath library

The math library supports C library mathematical functions. It supports number, angular conversion, trigonometric and hyperbolic functions. It also supports mathematical constants like π and e . You can use it with standard C functions. In the following example, we are importing the math library to use the `sin()` function in order to carry out a mathematical operation:

```
>>> import math
>>> math.sin (math.pi / 7.0)
```

Output:
0.43388373911755812

The `cmath` library is used to do complex mathematical



operations. The math and cmath library functions declaration is similar. The math library supports only standard C functions to perform mathematical operations and the cmath library is usually used in complex operations. This is the difference between math and cmath functions. In the following example, the cmath library and the `sin()` function is used to get an output in the form of a complex number:

```
>>> import cmath
>>> cmath.sin (cmath.pi / 7.0)
```

Operations with Python

Output:
(0.43388373911755812+0j)

Now we will see how the math library functions can be used in numeric operations. Initially we will take the *modf()* function. This function returns the fractional and integer part of a specified number in the function. In the following example, the *modf()* function is demonstrated:

```
>>> import math
>>> math.modf (12567)
(0.0, 12567.0)
>>> math.modf (12567.890)
(0.88999999999941792, 12567.0)
>>> math.modf (-12567)
(0.0, -12567.0)
>>> math.modf (0.12567)
(0.12567, 0.0)
```

Logarithmic and power functions are important in complex mathematics, used for power and square root calculations. In the following example, *log()*, *pow()* and *sqrt()* functions are demonstrated:

```
>>> import cmath
>>> cmath.log (28)
(3.3322045101752038+0j)
>>> math.pow (28, 45)
1.3246813146595588e+065
>>> cmath.sqrt (28)
(5.2915026221291814+0j)
```

In this example, power function is used with the math library and the other two functions are used with cmath library. The cmath library does not support the *pow()* function.

It is possible to convert angles defined in radians to degrees and vice versa. The math library has *degrees()* and *radians()* functions. These functions are used for

conversion. The *degrees()* function can convert radians to degrees and the *radians()* function can be used to convert degrees to radians. In the following example, the degrees and radians functions are demonstrated:

```
>>> import math
>>> math.degrees (90)
5156.6201561774087
>>> math.radians (5156.6201561774087)
90.0
```

Mathematical notations

Mathematical notations are used to write mathematical equations and formulas. Pre-calculus symbols like '+' and '-' are widely used in many equations to solve problems. In Python you can use these notations very effectively for large calculations. In the following example, '+', '-' notations are used:

```
>>> a=150.50; b=80.45; c=12
>>> d = a+b-a+c+a+b
>>> print d
323.4
```

It also supports some advanced notations like '/', '*'. In the following example, these notations are used. Here you can solve inner and outer bracket equations. You can solve different equations in Python with the storage of equations as variables and use these variables in future calculations.

```
>>> a=12345678; b=-50; c=0.0078
>>> a/c + b*a
965495330.76923084
>>> (a/b + c/b) * (b*a + c*a)
152392260079127.25
>>> d = (a+b * c+a)
>>> e = (a-b * c-a)
>>> d + e
24691356.0
```

Absolute value expressions

Absolute value expressions are used to find absolute value of the given mathematical function or expression that is enclosed in closed parenthesis where the left parenthetical symbol is considered automatically present.

In Python you can simplify absolute value expressions easily. In the following example, the `abs()` function is used. It returns absolute values that are defined in the function.

```
>>> a=10; b=80; c=20
>>> d = abs ( a+b ) + abs (b+c)
>>> e = abs ( a-b ) + abs (b-c)
>>> abs ( d+e )
320
>>> a = 10.75; b=80.25; c=20.25
>>> d = abs ( a+b ) + abs (b+c)
>>> e = abs ( a-b ) + abs (b-c)
>>> abs ( d+e )
321.0
```

Python complex number arithmetic

Python supports complex number arithmetic. In the following example snippet, we have used the square-root functions, which are carried out for different numbers. You can make the `ljsqrt()` function to evaluate these arithmetic operations. Note that it will not work for negative numbers. In this case you need to use complex numbers in the form of $(x + j)$ where x is a number. For example, when you want to calculate the square root of 9, then you can directly use the `ljsqrt(9)` function. This will not work for `ljsqrt(-9)`. You have to use the $(-9 + 0j)$ number to calculate the square root of the -9 number.

```
>>> def ljsqrt(x):
...     return x**.5
...
>>> ljsqrt(9)
3.0
>>> ljsqrt(-9+0j)
(1.8369095307335659e-016+3j)
>>> ljsqrt(-9)
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
  File "<stdin>", line 2, in ljsqrt
ValueError: negative number cannot be raised to a fractional
power
```

You can add, subtract complex number equations as shown in the next example. Here, addition and subtraction is used to get the result from different equations. Absolute value expression is used to get the absolute final value.

```
>>> a = (10+0j); b= (-20+0j); c= (0.15+0j)
>>> d= (a+b) + (c+b)
>>> e= (a-b) + (c-b)
>>> d + e
(20.299999999999997+0j)
```

```
>>> abs ( d + e )
20.299999999999997
```

Simplifying and solving

In algebra, simplifying and solving expressions have some rules. You can use these rules in expressions for normal evaluation with parentheses, and for mathematical operations like addition or subtraction. Note that while using the associative and distributive property of algebra you should apply the same rules for both sides of the equations where the equal sign (`=`) is used to differentiate these two equations. In this section you will learn how to simplify equations with Python. When you assign two equations with the `==` sign, then it will give a true or false value depending upon the equations' values.

In the following example, two expressions are given. It will return 'True' if these two equations are equal; otherwise it will return 'False'.

```
>>> a = 80; b = 20; c = 80; d = 20
>>> a + b
100
>>> c + d
100
>>> (a+b) == (c+d)
True
>>> (a+b) == (c-d)
False
```

Any mathematical equation can be defined as a mathematical representation of two expressions that are separated by an equal sign (`=`). It is considered that the left side of the expression has the same value as the right side of the expression. These expressions may contain variables to denote some values in the equation. Solving of mathematical equations involves two phases: manipulation of the equation and finding the values of variables declared in the expression. In algebra you can write equations as $x+10 = 100$ where you need to find the actual value of x from the given expression. Here you try an add or subtract operation for both sides to obtain the actual value of x .

Let us try this out for the equation $x+10 = 100$. You have to subtract 10 from 100 to obtain the value of x . In the following example snippet, we have demonstrated how to solve equations. Here, x and y are defined with a specific value. The first part of the code in this example shows how to find values from solving complex equations. The second part of the code uses simplifying rules and gives results in `True/False` terms. Here, both sides of the equations are unequal so it shows up a 'False' result. In this way we can solve complex expressions where the equations contain predefined variables.

```
>>> x = 12.45
>>> 78 * x + (x + 10)
993.54999999999995
```



```
>>> (x + 10) * (x - 10) + (x*2 + 234)
313.90249999999997
>>> (x + y) + (x/y) == (x - y) - (y/x)
False
```

NumPy and Python

NumPy is useful for scientific computing. It supports multi-dimensional array objects. It is capable of performing linear algebra functions and Fourier transforms. It can be embedded with FORTRAN and C/C++ code. NumPy is used as an alternative to MATLAB. NumPy's advantage over MATLAB is that it is open source and a complete object oriented language. It was written by Jim Hugunin in 1995 and is based on the Numeric package, which is outdated. NumPy is a combination of the Numeric and Numarray package. Numeric supports cover array arguments and scalar operations. Numarray is faster for large arrays and a little slower for small arrays.

Let us try out the features of the NumPy package. For this you need to import the NumPy library—the following example snippet shows how. The *array()* function is used to store array elements. In this example, the addition, subtraction and multiplication of two matrices have been shown:

```
>>> from numpy import *
>>> l = array ((10.25, 20.50, 30.75))
>>> m = array ((70.75, 80.25, 90.50))
>>> n = l + m
>>> print n
[ 81.0 100.75 121.25]
>>> n = l - m
>>> print n
[ -60.50 -59.75 -59.75]
>>> n = l * m
>>> print n
[ 725.1875 1645.125 2782.875 ]
```

Now try to find the matrix transpose with the use of the 'T' prefix to the matrix. For this you need to import the *matrix* and *linalg* library from NumPy. In the following example, finding the matrix transpose is demonstrated:

```
>>> from numpy import linalg
>>> from numpy import matrix
>>> a = matrix ([ [10, 20, 30], [40, 50, 60], [70, 80, 90] ])
>>> print a.T
[[10 40 70]
 [20 50 80]
 [30 60 90]]
```

It is possible to construct 2D arrays from equal-sized vector arrays. Python supports this functionality with the *hstack* function. If *m* and *n* are two vectors of the same length, then you can try *mn = [m, n]* to *hstack*. This is shown in the following example. Here you will get

the output as 2D array elements.

```
>>> m = array (0, 10, 4)
>>> n = array (4)
>>> mn = hstack ([m, n])
>>> print mn
[0 4 8 0 1 2 3]
```

A SciPy tutorial


SciPy is open source software. It is used in mathematics, science and engineering application development. It is pronounced as 'Sigh Pie'. The SciPy library depends on the NumPy library. It supports N-dimensional array manipulation. It is used for image and signal processing, genetic algorithms and as an ODE solver.

You can use it with the NumPy library to perform complex operations. In the following example, SciPy is demonstrated to find out whether the matrix is regular and to find the inverse of a matrix.

```
>>> from numpy import matrix
>>> from scipy.linalg import det
>>> from scipy.linalg import inv
>>> x = matrix ([ [10, 20.56, 30.16], [40.56, 50, 60.25], [70.35, 80.90, 50.70] ])
>>> y = matrix ([15.25, 20.25, 10.75]).transpose ()
>>> print det(x)
14350.01812
>>> print inv(x)*y
[[-0.71549026]
 [ 0.50544609]
 [ 0.39830673]]
```

It is possible to calculate the integral of a given function with the *integrate* module available in SciPy. In the following example, the *quad()* library is used to evaluate the value of 'a'. In-built functions like Lambda can be used very efficiently. Lambda is an operator used in LISP and Python programming language to create closures in expressions.

```
>>> from scipy.integrate import quad
>>> quad (lambda a: a**5, 10, 100)
(166666500000.00003, 0.001850369857336887)
```

With the NumPy and SciPy library you can perform complex scientific and mathematical operations easily. Enjoy the richness of Python in scientific computing! 

REFERENCES:

- www.python.org
- www.onlamp.com

By: Suhas A. Desai works at Tech Mahindra. He writes on open source and security. In his free time, he volunteers for social causes.



A Security-Aware Nameserver



It's time to go beyond Bind and get started with MaraDNS.

To begin with, let's try to understand the recent DNS vulnerability that affected the Internet worldwide.

We all know that a DNS (domain name server) is used to obtain the IP address of a website from the website name. It is possible for malicious people to spoof the correct IP address of the website by manipulating the reply from the DNS server and inserting their own. So, users are directed to a wrong website, which could be malicious in nature and thus open avenues for the bad guys to collect information from them, without even resorting to phishing techniques!

It is not so easy, but neither is it too difficult either. There are 65,536 ports that can be used by the DNS server to send its replies back to the user. So, the bad guy has 1/65536 of a chance of getting it right. Though this is a very low probability, it is still realistically possible to be able to exploit it.

It has been known that malicious folk have spoofed data ready and there are scores of them who work together to share the spoils. But, what has been newly discovered is based on the following:

- The queried DNS server may not know the IP address of the given website;
- Thus, the DNS server forwards the request to another DNS server;
- The baddies may not be able to provide the IP address for *www.exampledomain.com*, but they can give out the IP address for the sub-domain, *sub.exampledomain.com*, by spoofing it with their malicious replies.

The chances of being able to get this right are very low, but the baddies can carry out the exercise as many times as they want. So, if they stumble upon the right port, they get lucky and are able to carry out their malicious acts.

There have been such attempts made in the past by trying to guess the secret number that decides the port to be used or trying to insert a new secret number to be accepted. But, they never worked. The new attack does not try to guess or make the DNS server accept the secret number; rather, it tries to guess which ports are being used so that the next IP address request coming in can be spoofed, as the secret number generated isn't random enough.

This is where DNS servers like MaraDNS stand out. They do return a truly random secret number. So now, the baddies are back to the old problem of trying to play catch up, as they cannot guess the port to be used.

What's MaraDNS?

MaraDNS is a security aware, lightweight DNS, which has the following advantages:

- *Security:* MaraDNS has a security history as good as or better than any other DNS server. It was not even affected by the recently-discovered DNS vulnerability mentioned earlier.
- *Well-supported:* MaraDNS has a long history of being maintained and updated. The most recent release was on August 4, 2008.
- *Easy to use:* A basic recursive configuration needs only a single three-line

configuration file. MaraDNS is fully documented, with both easy-to-follow tutorials, and a complete and up-to-date reference manual.

- *A small size:* It is well suited for embedded applications and other environments, where the server must use the absolute minimum number of resources possible. MaraDNS' binary is smaller than that of any other currently maintained recursive DNS server.
- *Open Source:* It has a two-clause BSD licence that is almost identical to FreeBSD.

Installation

MaraDNS should be available for your distribution, and it is recommended you use binaries from your distribution to run it. However, if that's not the case, you can download the latest release from www.maradns.org/download.html. Once you have the file on your system, open a terminal and execute the following commands:

```
tar -xjvf ~/maradns-1.3.07.09.tar.bz2
cd ~/maradns-1.3.07.09
make
```

Now, your binaries have been created and you're ready to test it—you need to be the root user henceforth:

```
cp server/maradns /usr/local/sbin
cp tools/duende /usr/local/sbin
mkdir /etc/maradns
mkdir /etc/maradns/logger
touch /etc/mararc
```

Now open the */etc/mararc* file in a text editor and add the following lines:

```
ipv4_bind_addresses = "127.0.0.1"
chroot_dir = "/etc/maradns"
recursive_acl = "127.0.0.1"
```

...following which, start MaraDNS as shown below:

```
/usr/local/sbin/maradns &
```

If you are behind a NAT, such as a router, you will need to forward Port 53 UDP. And if you are behind a firewall, you need to allow MaraDNS to connect to the Internet.

Testing it is as simple as issuing the following command:

```
dig @127.0.0.1 example.com
```

If you get a reply, you're good to go.

Now, try running it in daemon mode:

```
killall maradns
mkdir /etc/maradns/logger
/usr/local/sbin/duende /usr/local/sbin/maradns
```

```
killall maradns
```

If all is fine till here, you can proceed to actually install MaraDNS:

```
cd /home/<username>/maradns-1.3.07.09
make install
```

Now, we need to add */usr/local/sbin* and */usr/local/bin* to our PATH variable by appending the following line to the */etc/bashrc* file:

```
export PATH=$PATH:/usr/local/sbin:/usr/local/bin
```

That's it; you're done! Your server is ready to run.

Basic configuration

The foremost configuration involves setting the bind address in */etc/mararc*.

If you want to use it for just the local PC, it has already been set up in the previous section. If you want it to listen on more IP addresses, add them as follows:

```
ipv4_bind_addresses = "address1, address2, address3,..."
```

For example:

```
ipv4_bind_addresses = "127.0.0.1, 192.168.0.1"
```

recursive_acl is the variable used to set the list of IP addresses/range of IP addresses that can access the DNS server:

```
recursive_acl = "address1, range1/netmask,..."
```


For example:

```
recursive_acl = "127.0.0.1, 192.168.0.0/24"
```

If you are on a slow network, then you can change the *timeout_seconds* variable:

```
timeout_seconds = "10"
```

What is needed now is a configuration file to start MaraDNS via the *init* system, which is out of the scope of this article as different distributions have different methods. If you installed MaraDNS using your package manager, you should have a file */etc/init.d/maradns* or */etc/rc.d/maradns*.

This is just the tip of the iceberg, as MaraDNS can do a lot more. To know about it, visit www.maradns.org/tutorial/tutorial.html 

By: Mehul Ved is a FOSS enthusiast interested in technology. He is fond of the command line and networking.



Go for SOA with Open Source

Open source is a proven business and technology model, and SOA is a contemporary, mature and successful enterprise architecture style. This article briefly introduces different open source product choices in the SOA stack, like ESBs, Rule Engines, Registries and Repositories.

SOA stands for Service Oriented Architecture. While functions, methods and classes are functional abstraction levels of earlier architectures and software paradigms, SOA takes the abstraction of functional logic to a higher level called 'service'. A service is module of logic that contains meaningful business functions. For example, 'credit card verification' can be considered a service. A service may contain multiple classes and components in itself, depending on the size and complexity of the logic contained in it. The most popular technological form of services are Web services. SOA promotes using services, and composing and orchestrating them to form business processes. Services are more meaningful abstractions to business problems/solutions. Hence, SOA bridges the gap between IT and business. Broadly speaking, SOA promises the following three benefits:

1. *Business agility:* If an enterprise's architecture is SOA-based, it supports business agility by aligning IT closely to business. When enterprise business processes change, such changes can be accommodated quickly by recomposing

the business processes from services, and modified business processes can be launched quickly. Similarly, when the enterprise wants to launch a new business process, it is possible to do it in an agile way, by simply composing the services to form a new business process. When two enterprises are involved in mergers or acquisitions, if both their IT infrastructure is SOA-based, then the IT part of the merger is a lot simpler and quicker, and helps business decision-making.

2. *Technology flexibility:* SOA and Web services are built around sets of standards like XML, SOAP, UDDI, HTTP, etc. And these support loose coupling and interoperability. From the integration perspective, it solves the problems of incompatibility.

3. *New business opportunities:* Like most other technologies, SOA opens windows of opportunity for new businesses. Like Internet mail and mobile technologies have generated new businesses, SOA provides new business opportunities for existing enterprises. For example, SaaS is one such business model based on SOA. This section provides a high-level briefing of

SOA, setting the context before moving to the next sections. This is not a complete briefing. Interested readers should go through relevant information widely available on the Net.

SOA reference architecture

Studying the SOA product stack from the SOA reference architecture point of view provides better comprehension. Figure 1 shows a simple SOA reference architecture and the functionality of each layer. (Note that this reference architecture is drawn for the purpose of 'product type' mapping with each layer, to help readers' comprehension of 'product type' for the products discussed in this article. However, there are more fully developed SOA reference architectures available in the market from SOA vendors, independent consulting firms and standard bodies. This reference architecture has been intentionally kept simple.)

The service consumer layer addresses the requirements of presentation, consolidation of information, mash-ups, etc. Often, this layer may be outside the boundaries of the service hosting network. Business process layers address process composition, choreography and orchestration. The service layer is where the functionality is exposed as a service, usually a Web service. The software infrastructure layer is where packaged applications and custom applications of traditional software engineering belong. The vertical layers represent relatively more non-functional requirements. The integration layer addresses the connectivity and integration of services to form business processes. The security layer addresses the security needs of services, business processes and of software infrastructure. The governance layer encompasses all the remaining layers and addresses compliance, regulations, and monitoring issues.

Figure 2 depicts the 'product category/type' mapping with respect to the SOA reference architecture in Figure 1. All important and popular product categories have been mentioned in the diagram. Each product type has been briefly discussed in the following sections.

Products in the SOA stack

Now, let us discuss some of the important and popular product categories, with information on the open source options available. Here are the SOA product categories discussed in this article:

- SOAP toolkits
- Application servers
- Enterprise service buses (ESBs)
- Registries and repositories
- BPEL engines
- Rules engines
- Business process modelling tools

With each product category, important technical features of the product are discussed. These technical features can be helpful in evaluating a specific product's maturity and fitness-for-use for a specific requirement at

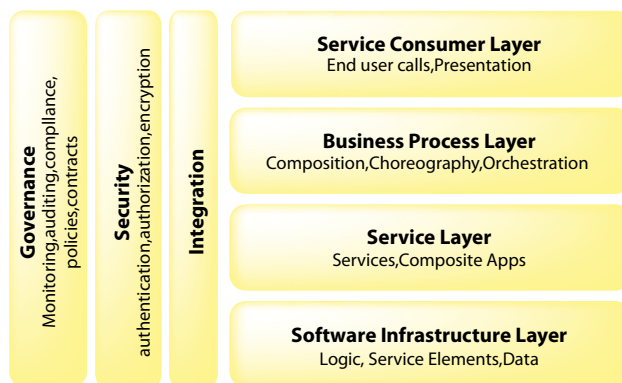


Figure 1: A simple SOA reference architecture

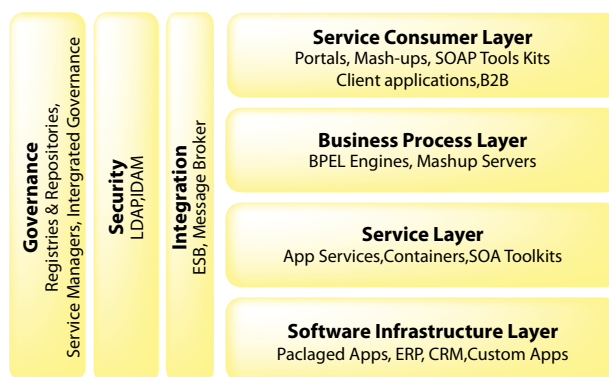


Figure 2: Product mapping with reference architecture

hand. Features mentioned here are not exhaustive.

SOAP toolkits

The SOAP toolkit is one of the preliminary product types for Web services-based SOA implementation. (Note that SOA can be realised with options other than Web services, but that aspect has not been discussed in this article.) In the Web services engineering, a SOAP toolkit may be an explicit product or may exist as a tightly-integrated component of an IDE, application server, etc. There are plenty of SOAP toolkits implementations available in the market, including open source and proprietary variants [2]. Apache Axis is one of the most popular open source choices in this product class. Apache Axis, CXF, gSOAP (C, C++ SOAP Toolkit) and PST! (PHP SOAP Toolkit), are some popular open source SOAP toolkits.

Some of the important technical features to look out for in a SOAP toolkit are:

- Support for SOAP 1.1
- Support for WS-2G specifications
- XML parser inside: SAX/DOM based
- Support for multiple transport protocols, viz., http, ftp, JMS, etc.
- Support for multiple programming languages—C++, Java, etc.
- Technology agnostic core modules of the toolkit
- Support to design WSDL graphically
- Support for generating WSDL from the service

- Support for generating client stubs or server skeletons from WSDL
- Support for RAD style interface for Web service development from popular programming languages like Java.

Apache Axis2 is considered as a 3G SOAP toolkit. For a list of open source Web services tools in Java, please refer to java-source.net/open-source/web-services-tools.

Application servers

While application servers are not specific to SOA, in this article the consideration is from the perspective of hosting Web services or their implementation components. Here, too, we have plenty of choice from open source. Apache Tomcat is one of the most popular and widely used among application servers that include Apache Geronimo, WSO2 WSAS, RedHat JBoss Application Server, ObjectWeb JonAS (an OW2 implementation [4]), GlassFish, Resin, Jetty, etc. Some of these servers host Java servlets only.

Some of the important technical features to look out in an application server are:

- Support for SOAP 1.1
- Support for multiple transport protocols, viz., http, ftp, JMS, etc
- Support for multiple programming languages—C++, Java, etc
- Security features
- Clustering, load balancing and failover
- OS support: availability on all popular OSs
- Enterprise service buses (ESBs)

ESB is an architectural pattern in the integration world. In contemporary technical literature, the word ESB is used to mean a software product category that implements the ESB pattern. This is a very important product that enterprises in their SOA journey usually implement in the early stages. Open source options are Apache Service Mix, Apache Synapse, RedHat JBoss ESB, Mule ESB, ObjectWeb Celtix, WSO2 ESB, IONA FUSE ESB, Chain Builder ESB from Bostech, jeeESB, etc. Some of these ESBs are JBI compliant and interoperable with other JBI compliant ESBs (open source or otherwise).

Some of the important technical features to look out for in an ESB are:

- Support for multiple data formats: text, XML, COBOL copy book, etc.
- Support for conversion across these different data formats (ESB may not hold the logic for data format conversion -- it might do so with the help of plug-ins).
- Mediation features
 - Transport protocol mediation across all popular protocols like http/https, JMS, ftp, SMTP, etc
 - Security assertions mediation across X.509, SAML, etc
- Security
 - Support for WS-Security
 - Support for X.509 and SAML
 - Support for SSL

- Routing
 - Support for content-based routing
 - Support for message correlation
- Governance features
 - Support for security, mediation, routing and other policy enforcement as a central PEP. (PEP, which stands for Policy Enforcement Point, is a location in software architecture where policy is enforced—for example, the service gateway.)
 - Monitoring and auditing support

Some of these features may be part of the ESB itself or be achievable through an external adapter that is available as a plug-in. For a quick summary of Java-based open source ESBs, refer to java-source.net/open-source/enterprise-service-bus.

Registries and repositories

The registry stores the indexes/metadata of resources and the repository stores the actual assets like business process elements, WSDLs, schemas, documents, etc. From the standards' perspective, a registry could be UDDI-based or ebXML-based. Registries and repositories host metadata and artefacts related to the services. In this article, registries and repositories are together referred as one product category. The open source choices are: Apache Service Registry, WSO2 Registry, SOA matrix's Registry, Mule's Galaxy, IONA Registry/Repository, Sun's Service Registry, freebXML RR, etc.

The registry helps the publication, discovery and consumption of the service lifecycle. The repository helps tracking of service as an asset, its relationship with other services, etc.

Some of the important technical features to look out in registries and repositories are:

- The organisation and management of resources
- The ability to embed itself inside an application
- Resource versioning
- A Web UI for viewing the artefacts/resources
- Support for resource dependency management
- Support for resource activity monitoring
- Support for resource life cycle management
- Support for WSDL validation wrt WS-I
- Support for standards like UDDI, ebXML
- Service publication features: interface, discoverability, metadata
- Collaboration features

For a comparison of two popular open source registries, please refer to wso2.org/library/3777.

BPEL/workflow engines

BPEL engines (a.k.a workflow engines) execute the business process in the form of BPEL and other process languages. Apache ODE, jBPM, OpenSymphony, OSWorkflow, ObjectWeb Bonita, ActiveBPEL, etc, are some open source BPEL engines.

Here are some of the important features to look for in

a BPEL engine:

- Support for WS-BPEL, BPEL4WS, etc
- Support for http WSDL binding
- Support for invocation of REST-style Web services
- Support for management and monitoring of business processes
- Availability of adapters to popular EISs
- Support for process languages other than BPEL (for example, jPDL), if required
- Support for popular application servers (for example, Apache Geronimo)
- Support for clustering
- Support for popular databases (for example, MySQL)

For a quick summary of Java-based open source workflow engine implementations, please refer to *java-source.net/open-source/workflow-engines*.

Rules engines

Business rules are very important elements of business processes. They provide a platform for developing, editing, testing and integrating business rules that are part of business processes. JBoss Drools, Open Lexicon, Apache OFBiz Rules Engine, Mandarax, SWEET Rules, OpenRules, JRuleEngine, Take, etc, are some of the open source rules engines available [6][7].

Here are some of the important features that a rules engine may have:

- Support for formats like XML, DRL, etc
- Support for decision tables
- Backward and forward inferencing
- Merging of rule bases/ontologies
- Rules storage in a database, with no database lock-in
- Availability as an Eclipse plug-in
- Intuitive UI to compose and edit rules
- Support for rule authoring, version control and management
- Rete algos
- Support for stateful and stateless execution modes
- Inline evaluation of rules
- Simple logical operations support (for example, AND, OR, EXISTS, etc)


Business process modelling tools

BPM (business process modelling) is an exclusive subject in itself; however, here it is discussed only with respect to SOA. (Note that SOA is a popular architectural style implementation choice for BPM.) Only BPM tools are discussed here, although some other BPM tools are loosely referred to as well. Business process modelling tools support modelling in a business process with the details of business process artefacts (a.k.a business items), resources, schedules and costs. These tools are very useful for business analysts, and provide a process simulation interface that would help analysts in the simulation of actual business processes. These simulations help understand the business process deeply, and provide the required insight to develop the to-be

process. Also, they provide support for automatic generation of code. Some of the most popular open source BPM tools are: Intalio Designer, Apache OFBiz, jBPM, OpenebXML, OSWorkflow, XFlow, etc. Many of these tools support more BPM functionality than mere business process modelling.

Some of the important technical features to look out for in a BPM tool are:

- Support for BPMN
- Support for BPML
- Business process code generation
- Support for process validation
- Support for business process simulation
- Availability as Eclipse plug-ins
- Intuitive user interface
- Graphical rules mapping
- Schema mapping editor
- XML and XPath support
- No database lock-in—support for multiple popular databases
- Support for integration with external LDAP servers for roles

This is a brief summary of open source SOA products, and not an exhaustive compilation. For example, SOA governance products, except for registries and repositories, are not covered. This article is meant to provide a brief overview of different open source options available in different SOA product categories. For more exhaustive details on these products and categories, readers need to explore the individual product's features and references given in this article. Also, note that some of these product categories are relatively loosely defined in the practice, and features of some product categories are also available in other product categories.  **END**

REFERENCES:

1. Open Source WebServices Tools in Java <http://java-source.net/open-source/web-services-tools>
2. A list of SOAP Toolkits available <http://www.soapware.org/directory/4/implementations>
3. List of SOA related products http://en.wikipedia.org/wiki/List_of_SOA_related_products
4. OW2 <http://www.ow2.org/view/About/OW2Consortium>
5. Quick summary of Java based ESBs <http://java-source.net/open-source/enterprise-service-bus>
6. Open Source Rules Engines in Java <http://java-source.net/open-source/rule-engines>
7. Open Source Rules Engines written in Java http://www.manageability.org/blog/stuff/rule_engines
8. Open Source Competition - Mule Galaxy vs. WSO2 Registry <http://wso2.org/library/3777>
9. Open Source Workflow Engines in Java <http://java-source.net/open-source/workflow-engines>

By: A S Rao. The author is a SOA technical architect and consultant, and can be reached at <asrao_tumi@yahoo.com>

A VOYAGE TO THE KERNEL



Part 6

Day Five—The End of the First Segment

We will now conclude the shell programming part of our voyage. In this column, I will try to review the tools described in earlier columns and apply those tools to solve little complicated problems --solutions that you may exploit while coding.

This part addresses three categories of readers: the first one is for newbies who have just started their experiments in shell; the next one is for intermediates and the last one for advanced users. Besides, I have skipped the illustration of some of the codes.

Let us fire up by trying a code for newbies:

```
#!/bin/bash

NUMBER=0

echo -n "Please enter number between 2 and 9"
read NUMBER

if ! [ $NUMBER -ge 2 -a $NUMBER -le 9 ] ; then
    echo "Please enter number between 2 and 9"
    exit 1
fi

clear

for (( i=1; i<=NUMBER; i++ ))
do
    for (( n=NUMBER; n>=i; n-- ))
    do
        echo -n " "
    done
    for (( j=1; j<=i; j++ ))
    do
```

```
        echo -n " $i"
    done
    echo ""
done
```

Figure 1 shows the execution of the code. For a detailed explanation of the mode of functioning, please refer to the earlier columns.

Now, let us write the code to find the reverse of a given number. (By this time, you must know why we use exit 1 in the code.)

```
#!/bin/bash

if [ $# -ne 1 ]
then
    echo "Usage: $0 number"
    echo "This will help you to find reverse of a number"
    exit 1
fi

number=$1
reverse=0
division=0

while [ $number -gt 0 ]
do
    division=`expr $number % 10`
    reverse=`expr $reverse \* 10 + $division`
    number=`expr $number / 10`
done

echo "Reverse number is $reverse"
```

Can you guess what the following code does? Else, give it a try and find out:


```
#!/bin/bash

echo "Enter number:"
read number
i=$number
while test $i != 0
do
    echo "$i"
    i=`expr $i - 1`
done
```

Sometimes, you may need to find out whether the user is logged in as a root user, especially when you write system tools (administration tools) in shell:

```
#!/bin/bash

ROOT_UID=0

if [ "$UID" -eq "$ROOT_UID" ]
then
    echo "Welcome, root."
else
    echo "Please login as root "
fi

exit 0
```

As the root user has \$UID 0, you can easily uncover this, by a conditional statement.

If you wish to have some fun in between, use the following code:

```
#!/bin/bash

echo "Enter number:"
read number

for (( i=1; i<=$number; i++ ))
do
    for (( j=1; j<=i; j++ ))
    do
        echo -n " |"
    done
    echo "_ "
done

echo "let's climb"
```

You can easily understand the above code by looking at *i* and *j* (and the increment factor associated with those). If you check the result of the operation "I" (and its alternative "_") you can guess what the final result will look like. A sample is

```
File Edit View Terminal Tabs Help

  1
 2 2
3 3 3
4 4 4 4
5 5 5 5 5
6 6 6 6 6 6
7 7 7 7 7 7 7
hacker@free-laptop:~$
```

Figure 1: Terminal output after execution of the first code for newbies

given in Figure 2.

Sometimes you may wish to extract some content from the Web. You have many methods to do that, using shell. I will show you how to fetch a Web page (say, an article from Wikipedia):

```
Enter number:
8
|_|_|_|_|_|_|_|
|_|_|_|_|_|_|
|_|_|_|_|_|
|_|_|_|_|
|_|_|_|
|_|_|
|_|
|_
let's climb
```

Figure 2: The "let's climb" code

```
#!/bin/bash

if [ -z "$1" ]
then echo "Usage: `basename $0` Wikipedia article name"
exit
fi

article=$1
URL='http://en.wikipedia.org/wiki/'
wget -O ${article} "${URL}${article}"
exit $?
```

A demo of the code is shown in the following snippet:

```
hacker@free-laptop:~$ /home/hacker/Desktop/a H
--12:07:16-- http://en.wikipedia.org/wiki/H
=> `H`

Resolving en.wikipedia.org... 208.80.152.2
Connecting to en.wikipedia.org|208.80.152.2|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 66,846 (65K) [text/html]

100%[=====] 66,846 2.14K/s ETA 00:00

12:08:00 (2.08KB/s) - `H` saved [66846/66846]
```

If you need to crawl through some special pages, you can add some suffix to the URL. For example, if you need the printable version, you can add something like:

```
suffix=&printable=yes
```

Then add *`\${suffix}* to *wget*.

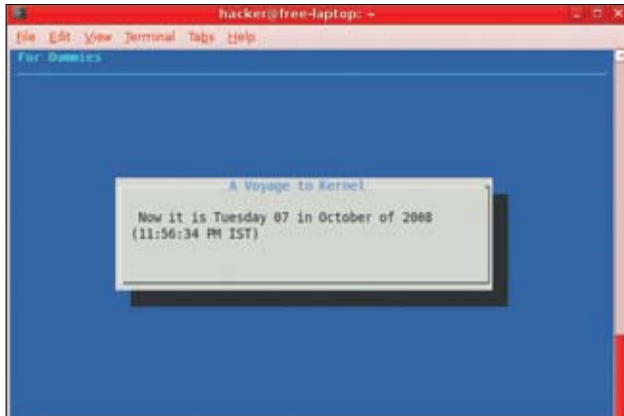


Figure 3: What's the system time?

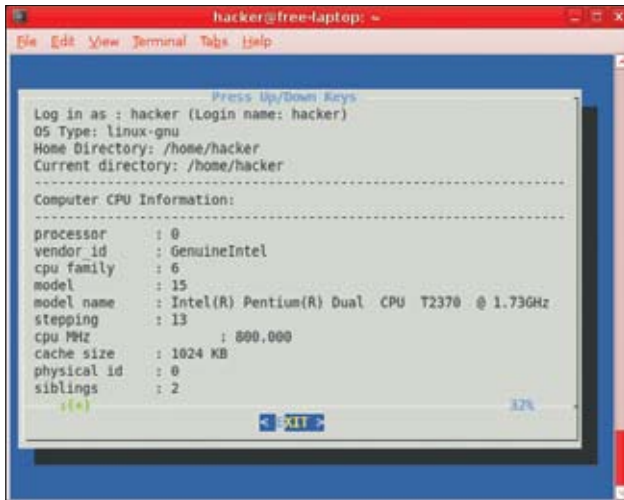


Figure 4: A tool that displays system information

Now, let's see how to capture the keystrokes using shell:

```
#!/bin/bash

keystrokes=10

old_pref=$(stty -g)

echo "Enter $keypresses keystrokes:"
stty -icanon -echo

pressed_keys=$(dd bs=1 count=$keystrokes 2> /dev/null)

stty "$old_pref"

echo "You pressed the \"$pressed_keys\" keys."

exit 0
```

You can see that we are able to disable the canonical mode and local echo in this. And *old_pref* is used to restore the old preference settings that are saved using *\$(stty -g)*.

You may also change the keystrokes value to your

preferred value.

While executing some shell-based commands, you may need to change the working directory automatically to the one in which the shell script is located. The following code does this:

```
#!/bin/sh

directory=`pwd`
for cmd in *
do
if test -d $directory/$cmd
then
cd $directory/$cmd
while echo "$cmd:~$"
do
read cmdmd
do
eval $cmdmd
done
cd ..
fi
done
```

If you wish to display the system time in your program, you can use the following code:

```
#!/bin/bash

temph=`date | cut -c12-13`
dat=`date +%A %d in %B of %Y (%r)`
dialog --backtitle "For \"
--title "A Voyage to Kernel\"
--infobox "\n Now it is $dat" 7 50
```

Figure 3 shows a demo. If you wish to have a dynamic greeting, you can use conditional statements and link them to the system timings.

Now, while writing admin tools you may have to display information regarding the system. The following code illustrates their usage (see Figure 4 for the demo):

```
#!/bin/bash

user=`who | wc -l`
echo -e "Log in as : $USER (Login name: $LOGNAME)" >> /tmp/info.
tmp.01.$$$
echo -e "OS Type: $OSTYPE" >> /tmp/info.tmp.01.$$$
echo -e "Home Directory: $HOME" >> /tmp/info.tmp.01.$$$
echo -e "Current directory: `pwd`" >> /tmp/info.tmp.01.$$$

echo -e "-----" >> /tmp/info.
tmp.01.$$$
echo -e "Computer CPU Information:" >> /tmp/info.tmp.01.$$$
echo -e "-----" >> /tmp/info.
tmp.01.$$$
cat /proc/cpuinfo >> /tmp/info.tmp.01.$$$
```

```
dialog --backtitle "A Voyage to Kernel" --title "Press Up/Down
Keys " --textbox /tmp/info.tmp.01.$$$ 21 70
```

```
rm -f /tmp/info.tmp.01.$$$
```

Similarly, you can display other information as well—say that related to computer memory, the hard disk, filesystem, etc.

```
#!/bin/bash
```

```
echo -e "-----" >> /tmp/info.tmp.01.$$$
echo -e "Computer Memory Info:" >> /tmp/info.tmp.01.$$$
echo -e "-----" >> /tmp/info.tmp.01.$$$
cat /proc/meminfo >> /tmp/info.tmp.01.$$$
echo -e "-----" >> /tmp/info.tmp.01.$$$
echo -e "Hard disk info:" >> /tmp/info.tmp.01.$$$
echo -e "-----" >> /tmp/info.tmp.01.$$$
echo -e "Model: `cat /proc/ide/hda/model` " >> /tmp/info.tmp.01.$$$
echo -e "Driver: `cat /proc/ide/hda/driver` " >> /tmp/info.tmp.01.$$$
echo -e "Cache size: `cat /proc/ide/hda/cache` " >> /tmp/info.
tmp.01.$$$
```

```
echo -e "-----" >> /tmp/info.tmp.01.$$$
echo -e "File System :" >> /tmp/info.tmp.01.$$$
echo -e "-----" >> /tmp/info.tmp.01.$$$
cat /proc/mounts >> /tmp/info.tmp.01.$$$
```

```
dialog --backtitle "A Voyage to Kernel" --title "Press Up/Down Keys
" --textbox /tmp/info.tmp.01.$$$ 21 70
```


```
rm -f /tmp/info.tmp.01.$$$
```

The next segment

I was planning to take a leap into kernel programming directly. But from the e-mails that I have received, I understand that many readers are new to areas like writing device drivers. And some readers are unfamiliar with tasks like kernel compilation. Considering the requests from beginners and intermediates, I am changing our voyage schedule.

Instead of going directly into kernel programming, I shall introduce you to a new segment dealing with the mathematical skills required for problem solving. I am of the outlook that computer science has got nothing to do with computers. It is the science of problem solving using algorithms.

Even in kernel programming, you can use many of these tools. This will enable intermediates to acquire more mathematical skills in programming, which are indispensable when playing around with the kernel. But I have ensured the layout of the next segment suits all programmers. Hence, even if you don't wish to meddle much with the kernel, you will find these tips useful for writing all types of algorithms.

Stay tuned!  **END**

VOYAGE MUSIC

Let's end our voyage by playing some music! Our journey to this locale is about to draw to a close and this piece further illustrates the clout of shell. You may see a reference to `/dev/dsp`, which is actually the Digital Signal Processor. You can also vary the tune, sound, etc, for better results. If you are good in music, you will find that I have used the European notation in the code. Now let's play the notes:

```
#!/bin/bash
```

```
duration=1000
volume=${'xff'} # Max volume = \xff
mute=${'x80'} # No volume = \x80

function voyage_music () # Voyage music note Hz in bytes
{
    for t in `seq 0 $duration`
    do
        test $(( $t % $1 )) = 0 && echo -n $volume || echo -n $mute
        done
    }
}
```

```
e=`voyage_music 50`
g=`voyage_music 42`
a=`voyage_music 39`
b=`voyage_music 40`
c=`voyage_music 21`
cis=`voyage_music 23`
d=`voyage_music 21`
e2=`voyage_music 22`
n=`voyage_music 32767`
```

```
echo -n "$g$e2$d$c$d$c$a$g$n$g$e$n$g$e2$d$c$c$b$c$cis$n$cis$d \
n$g$e2$d$c$d$c$a$g$n$g$e$n$g$a$d$c$b$a$b$c" > /dev/dsp
```

```
exit ?
```



By: Aasis Vinayak PG. The author is a hacker and a free software activist who does programming in the open source domain. He is the developer and CEO of the Mozhi Search engine. His research work/publications are available at www.aasisvinayak.com

CodeSport

Welcome to another instalment of CodeSport. In this month's column, we discuss a method of algorithm complexity analysis known as amortised cost analysis.

Thanks to all the readers who sent in their solutions/comments to the problems we discussed in last month's column. The takeaway problem was a variant of Dijkstra's single source shortest path (SSSP) algorithm, namely that of finding the shortest paths in directed acyclic graphs (DAG) with no negative weighted edges. The challenge was to come up with an algorithm that could solve the SSSP problem for DAGs in $O(V+E)$ time complexity.

First of all, note that shortest paths are well defined for directed acyclic graphs, even if the graph has negative weight edges. This is because in a DAG, no negative weight cycles can exist. Remember that Dijkstra's SSSP algorithm repeatedly selects a vertex that has the shortest path estimate from the source and relaxes the edges leaving that vertex. However, in case of DAGs, the relaxation order is quite straightforward. We relax the edges of the DAG according to a topological sort of the DAG's vertices. We have discussed topological sorts in one of our earlier columns. Recall that a topological sort of a DAG is a linear ordering of its vertices such that if there is an edge from vertex 'u' to vertex 'v' in G, then 'u' appears before 'v' in the topological ordering of G. A valid topological ordering is possible only for directed acyclic graphs. By processing the vertices in topological order, we ensure that we process all vertices that have directed edges to 'v' before processing 'v'. Here is the pseudo-code for SSSP on a DAG:

```
// w is the edge weight matrix of DAG G
// s is the source
Shortest_path_on_DAG(G, w, s)
{
    Topologically sort the vertices of G

    //d[v] is the shortest path estimate for v from s
    For (vertex v = 1 to n)
    {
        d[v] = infinity;
    }
    d[s] = 0;

    For each vertex u taken in topologically sorted order
    {
        For each vertex v adjacent to 'u'
            Relax(u, v, w);
    }
}
```

We know that a topological sort has a complexity of $O(V+E)$.

The second 'for' loop where we process vertices in topological order executes V iterations. For each vertex, we examine the edges that leave that vertex exactly once and relax that edge. Each relaxation step takes a constant time. Hence the complexity of our solution is $O(V+E)$.

Typically, Dijkstra's SSSP algorithm is implemented by using a binary heap to maintain the shortest path estimates from source for all vertices. The algorithm uses an 'Extract-min' operation on the binary heap to select the vertex with the shortest path estimate from the source at each iteration. By using a binary heap, we can perform *extract_minimum* in $O(\log V)$. Hence, the total time comes down to $O(V \log V + E \log V)$. Is it possible to reduce the complexity further? It is possible to do this by using an advanced data structure known as Fibonacci heaps. This data structure reduces the amortised time complexity of Dijkstra's SSSP algorithm to $O(V \log V + E)$.

In this month's column, we will start off by looking at the concept of amortised analysis and then discuss Fibonacci heaps in the subsequent columns.

When we analysed the complexity of algorithms, so far, we have seen the best case time complexity, the worst case time complexity and the average case time complexity based on the inputs to the algorithm. Remember that average case analysis is based on the running time over all possible inputs. Amortised analysis is a form of average analysis, but here we are not considering the average running time for a single operation over all possible inputs, but the amortised analysis is by taking the average over N consecutive operations. Amortisation is a clever trick used by accountants to average out large one-time costs over long periods of time. For instance, one of the common examples of amortisation is the case of calculating uniform payments for a loan, even though the borrower is paying interest on less and less capital over time.

Let us consider the example of incrementing a binary counter. Given an 'n' bit binary counter, here is the pseudo code for incrementing a binary counter B:

```
Increment (B)
{
    int i = 0;

    while (B[i] != 0)
    {
        B[i] = 0;
        i = i + 1;
    }
    B[i] = 1;
}
```


What is the running time of the increment algorithm?

The running time depends on the current configuration of the counter passed as the input to 'increment'. If the first 'k' bits are ones, then the increment takes $O(k)$ time. Hence, if the number passed as input to increment algorithm is between 0 and N , then 'increment' takes $O(\log N)$ in the worst case time since the binary representation for 'n' is $\text{floor}(\log N) + 1$ bits long.

Let's suppose we call the increment consecutively N times, starting with the counter at 0. Since the worst case running time of the increment is $O(\log N)$, we arrive at the $O(N \log N)$ worst case time complexity for the N consecutive increments of the counter. Although the upper bound of $O(N \log N)$ is correct, it is not the tight bound. We can show that the N consecutive increments starting from the counter at zero, take only $O(N)$ time complexity.

As mentioned before, if the first 'k' bits are 1, then the increment takes $O(k)$ times since all the 'k' bits need to be flipped when we increment the counter. Now, if we write down the binary counter values starting from 0 to $N-1$, we can see that each bit in the binary representation does not flip every time we increment. We can see that the least significant bit denoted as $B[0]$ flips every time, and the next significant bit $B[1]$ flips only every other time (for example, from 0001 to 0010 or 0011 to 0100). $B[2]$ flips every fourth time and, in general, $B[i]$ flips every 2^i th time. If we start from the counter at 0, each bit flips exactly $N/2^i$ times when we call the increment for N consecutive times. Hence, the total number of bit flips for the N consecutive increments is given by $N/2^i$ for $0 \leq i \leq \log N$. This can be approximated by $N/2^i$ for $0 \leq i$ and hence is equal to $2n$. Taking the average of over N operations, we see that each increment flips only 2 bits and hence runs in constant time. Thus, the amortised cost of the increment is constant, wherein the worst case cost of the increment is $O(\log N)$.

There are a number of different methods for deriving amortised cost bounds for a sequence of N operations. Three commonly used methods are the aggregate method, the accounting method and the potential method. The amortised cost analysis we discussed for the binary counter increment is based on the aggregate method. In the aggregate method of amortised analysis, we find the worst case cost of a sequence of N operations as $T(N)$ and then calculate the amortised cost of the single operation as $T(N)/N$.

Accounting methods are different from aggregate methods in that we assign charges to different operations, with some operations assigned charges more or less than what they actually cost. We charge more for certain operations so that the saved charge can be used up when an operation that has a high actual cost, occurs. Let us consider the example of our binary counter increment. When we change a bit from 0 to 1, we charge the operation a cost of 2 units. One unit of cost is for performing the actual flip from 0 to 1. The other unit of cost can be considered as an extra credit that is stored with that bit and will be used to pay for the cost of resetting the same bit to zero. Consider the case when the counter is in state 0011 and the increment is called on it. Now when bits $B[0]$ and $B[1]$ were flipped from 0 to 1, they were charged 2 units each. Hence, now each of them has an extra unit of credit stored with them that can be used for resetting them to 0 in the current call to increment. Remember that we associated a cost of $O(k)$ with the increment if the first 'k'

bits were ones, since all of the 'k' bits needed to be reset to 0 during this increment. However, now each of those 'k' bits have an extra unit of credit (which was charged when they flipped from 0 to 1) to pay for this transition from 1 to 0. This ensures that each call to increment incurs only a constant cost as we use up the previous credit when the bits need to be reset from 1 to 0.

Note that in accounting methods, we can charge an operation a cost different from its actual cost. The amount we charge an operation is called an amortised cost. For example, when a bit is flipped from 0 to 1, we charge it a cost of 2, whereas its actual cost is 1 only. When the amortised cost we charge an operation exceeds its actual cost, the difference is stored as credit with a part of the data structure. This credit can be used later to pay for operations whose actual cost is more than their amortised cost. In the case of a binary counter, we associate the credit with each bit, and this extra credit is used to pay for the resetting of that bit from 1 to 0. Note that the amortised cost analysis must ensure that the total amortised cost on the sequence of N operations must be an upper bound on the total actual cost of the sequence of N operations. Only then we can use the amortised cost analysis to show that in the worst case, the average cost per operation is small. Hence, we have the inequality:

$$\sum_{i=1}^n \text{amortized cost } c_i \geq \sum_{i=1}^n \text{Actual cost } a_i$$

We will not discuss the potential method of amortised analysis here except to state that the potential method associates prepaid work with the whole data structure as a potential rather than with individual pieces of data structure, as in the "accounting method". The potential method is the most rigorous method of amortised analysis and hence is the hardest. Interested readers can find details on potential amortised cost analysis in "Introduction to Algorithms" by Carmen, Leiserson and Rivest.

For this month's takeaway problem, let us apply amortised analysis to a data structure that can support the following two operations on a set of integers, namely:

- (a) `Insert(S,x)` inserts x into the set S
- (b) `delete_half_set` deletes $\text{ceil}(S/2)$ elements from the set S .

The data structure you design should be such that a sequence of M operations runs in $O(M)$ time. In other words, the amortised complexity of both `insert` and `delete_half_set` should be constant, when a sequence of M operations consisting of `insert` and `delete_half_set` are performed on the data structure.

If you have any favourite programming puzzles that you would like to discuss at this forum, please send them to me. Feel free to send your solutions and feedback to sandyasm_AT_yahoo_DOT_com. Till we meet again next month, happy programming! 

Sandya Mannarswamy is a specialist in compiler optimisation and works at Hewlett-Packard India. She has a number of publications and patents to her credit, and her areas of interest include virtualisation technologies and software development tools.



S.G. GANESH

What Went Wrong?

C programming is fraught with traps and pitfalls. In this column, we'll see some deceptively simple programs or expressions, which have bugs. Turn into a 'bug detective' and find out what went wrong!

- 1 The conditions `(i % 2 == 1)` and `(i & 0x1 != 0)` fail to check if `i` is an odd number or not. Why?
- 2 We want to resize an image by 'scale' factor; for that we want to make sure that scale is not equal to 0.0 or 1.0 (it does not make sense to scale an image by 0.0, which is not possible, or by 1.0, which will not have any effect on the image size). So, the programmer wrote a code like this:

```
if( (scale != 0.0) || (scale != 1.0) )
    imageSize /= scale;
```

What is wrong with this solution?

- 3 An experienced programmer wrote this simple program, but he was surprised that his program core-dumped instead of greeting "Hello world!".

```
// File: c.c
const char greetings[20] = "Hello world!";

// File: d.c
#include <stdio.h>
extern const char *greetings;
int main() {
    printf(greetings);
}

// in command-line, it is compiled and run as follows
// $ cc c.c d.c
// $ ./a.out
// core dump message here ...
```

Can you find what went wrong?

- 4 This simple factorial program did not work correctly!

```
int fact(int x) {
    if(x <= 1)
        return x;
    else
        return (x * fact(--x));
}
```

Can you find what went wrong?

- 5 You wanted to write a small code that will keep printing "processing... hold on.." but nothing got printed!


```
int main() {
    for(;;) {
        printf("processing... hold on...");
        sleep(1); // sleep is in <unistd.h>
    }
}
```

Can you find what went wrong?

- 1 The first expression fails if `i` is a negative number. The second expression is read as `(i & (0x1 != 0))` since explicit parenthesis is missing! It's better to use the check `((i % 2) != 0)`.
- 2 This `if` condition can result in divide-by-zero error if scale is 0.0! The condition should use `&&` instead of `||`, as in `((scale != 0.0) && (scale != 1.0))`! The programmer, while reading the statement (say from requirements document) "scale is not equal to 0.0 or 1.0" might have directly (and wrongly) translated it into code to use `||` operator; in fact, it should be `&&` operator.
- 3 The declarations `const char greetings[20];` and `const char *message;` are not equivalent! The extern declaration of the message should be `extern const char greetings[20];`.
- 4 The problem is the side-effect pre-decrement operator, so replace it as `(x - 1)`, as in:

```
return (x * fact(x - 1));
```

Since the value of `x` doesn't get changed in the expression itself, and a temporary variable is used for passing the result of `(x - 1)` to `fact` function, it works well.

- 5 The sleep function suspends the process from execution (for the specified number of times), but as it's an infinite loop, the program might never get to flush the stream and hence you might not see any output printed from the program. 

S.G. Ganesh is a research engineer in Siemens (Corporate Technology). His latest book is "60 Tips on Object Oriented Programming", published by Tata McGraw-Hill in December last year. You can reach him at sgganesh@gmail.com

LINUX JOBS

Post: Sr. Oracle DBA
Company: Sagarsoft (India) Ltd.
Profile: Should have technical knowledge of Oracle versions 9i & 10g, 8i, 11g desirable etc.
Exp.: 5-10
Location: Bangalore
Email: careers@sagarsoft.in

Post: QA Analyst
Company: Jean Martin
Profile: Ideal candidate should have good understanding of FIX protocol, good UNIX/ LINUX skill, domain Knowledge of US Financial Markets etc.
Exp.: 4-9
Location: Chennai
Email: umab@jeanmartin.com

Post: Solution Analyst
Company: Arcot R & D Software Pvt. Ltd.
Profile: Should have hands on experience of complete Software Development Life Cycle process & designing solutions.
Exp.: 7-9
Location: Bangalore
Email: sanjana@arcot.com

Post: Project Lead- IN, VAS
Company: Megasoftware Ltd.
Profile: Experience from vendor side would be an added advantage.
Exp.: 6-8
Location: Hyderabad
Email: rakeshn@xius-bcgi.com

Post: Platform Engineer- UNIX/ Linux kernel, RISC processors
Company: Careernet Technologies Pvt. Ltd.
Profile: Applicant should possess excellent knowledge of UNIX/ Linux kernel with understanding of X78 or RISC processors, device drivers & IO modules in kernel.
Exp.: 5-10
Location: Bangalore
Email: karthik.v@careernet.co.in

Post: JCAPs Technical
Company: Pantaloon Retail India Ltd.
Profile: Incumbent should have experience in Sun SeeBeyond JCAPs or ICAN. Should be strong at JCAPs installation, development of JCDs, OTDs, configuring JCAPs eWays etc.
Exp.: 7-10
Location: Ahmedabad
Email: ritesh.devadiga@pantaloon.com

Post: Component Developer
Company: Photon Infotech Pvt. Ltd.
Profile: Should have knowledge of J2EE, .Net, Linux, UNIX, Windows 2003 server etc. Familiarity with Routers, Switches, Firewalls configurations is also required.
Exp.: 6-11
Location: Bangalore
Email: prabhu@photoninfotech.com

Post: Oracle DBA
Company: Omnitech Infosolutions Ltd.
Profile: Hands on Oracle Export & Import is required with exposure to manual recovery & backup techniques, RMAN backup etc.
Exp.: 5-10
Location: Mumbai
Email: alan@omnitechindia.com

Post: SQL Development Professional
Company: Agilis Information Technologies Intl. Pvt. Ltd.
Profile: The candidate should be good in PL/ SQL development. Should have hand on experience on coding related to oracle. Exposure to Unix, Linux is mandatory.
Exp.: 4-9
Location: Delhi
Email: careers@agilisinternational.com

Post: Project Manager- Operations
Company: CollabNet Software Pvt. Ltd.
Profile: Understanding of networking, data centers, software development process, customer support practices etc.
Exp.: 10-12
Location: Chennai
Email: recruitment@collab.net

Post: Testing Lead
Company: Agilis Information Technologies Intl. Pvt. Ltd.
Profile: Should have experience in testing of products or modules in Test Lead Role, Querying the Oracle database, long or short term programming (Automation) etc.
Exp.: 4-7
Location: Delhi
Email: careers@agilisinternational.com

Post: Sr. Unix Administrator
Company: Atlantium Info Services Pvt. Ltd.
Profile: Candidate should have 5+ yrs of experience in the same domain.
Exp.: 5-9
Location: Bangalore
Email: sowmyarai@atlantium.net

Post: Build & Release Engineer
Company: Tanla Solutions Ltd.
Profile: Must Have experience of SCM (Software Control Management), preferably with open source CVS & subversion experience with Unix OS with emphasis in Linux environment.
Exp.: 4-7
Location: Hyderabad
Email: sharmila.ankireddy@tanlasolutions.com

Post: Database Administrator
Company: FCS Software Solutions
Profile: Must know usage of different partitions like hash, list & range, with strong knowledge of external tables & data loading activities.
Exp.: 4-8, Location: Pune
Email: mrjkkumar@fcsitd.com

Post: Performance Tuner
Company: Mastek Ltd.
Profile: Incumbent should have experience in performance tuning. Should be technically sound in Oracle & UNIX /Linux technology. Knowledge of DBA activities is also required.
Exp.: 7-9, Location: Mumbai
Email: deepitg@mastek.com

Post: Sr. J2EE Developer
Company: BPA Technologies Inc.
Profile: Should have software development experience using standard web technologies with strong understanding of version control tools like subversion or CVS.
Exp.: 5-7
Location: Visakhapatnam
Email: seshu@bpatech.com

Post: Sr. Java
Company: BPA Technologies Inc.
Profile: Should have software development experience using standard web technologies. Must have strong relational database experience, Linux/ Unix experience etc.
Exp.: 5-8
Location: Visakhapatnam
Email: seshu@bpatech.com

Post: Webmethods Integration Developer
Company: Verion Technology Solutions Pvt. Ltd.
Profile: Applicant should have 5+ yrs of experience in the same domain.
Exp.: 5-7
Location: Bangalore
Email: kaladhar.vemuri@verion.com

Post: Oracle DBA
Company: Akshay Software Technologies Ltd.
Profile: Candidate should possess knowledge of shell scripts, oracle 9i/10g, OCP certification, physical & logical standby database configuration etc.
Exp.: 4-5, Location: Mumbai
Email: hrd@akshay.com

Post: Informatica- SSE
Company: Damco Solutions Pvt. Ltd.
Profile: Incumbent should have design, development & testing experience in some data warehousing project involving EDW/ Star Schemas with excellent SQL & PL/ SQL skills. Familiarity with Core Java, OOP skill working knowledge of Linux/Unix is also required.
Exp.: 5-9
Location: Bangalore
Email: mukeshk1@damcogroup.com

Post: Sr. Software Test Engineer
Company: IC Infotech
Profile: Should have exposure to Load Runner, to LINUX environment, testing methodology.
Exp.: 4-6
Location: Delhi/NCR
Email: kanchan.g@c2il.com

Post: Sr. Unix Administrator
Company: Yahoo Software Development India Pvt. Ltd.
Profile: Candidate should have good exposure to Perl/ Shell scripting.
Exp.: 7-12, Location: Bangalore
Email: barrgavi@yahoo-inc.com

Post: Sr. System Administrator
Company: Renaissance Support Services Pvt. Ltd.
Profile: VMW experience is highly desirable. Experience with Unix/ Linux operating systems, NAS and/or SAN storage & databases will be an added advantage.
Exp.: 7-12
Location: Bangalore
Email: hr@renaissances.com



Get the Best Paying Jobs @ naukri.com
2,00,000 Jobs to choose from!

35,000 Recruiters | Get Jobs by email | Block Your Current Employers

Log on to www.naukri.com

Powered by
naukri.com
India's No. 1 Job Site



Using your Windows fonts

Although many users probably dual boot their Linux systems with Windows, the font server isn't configured to use the fonts you have on your Windows partition. It's an easy fix, so you might as well put them to use! (This tip probably applies to other rpm-based distributions as well, apart from Mandriva.)

First, make sure your Windows partition is mounted automatically on boot. Type `ls /mnt/windows` (or wherever your Windows C partition is mounted). If there are files there, you are fine. If not, start the Mandrake Control Centre and run the disk partition program. Click on the Windows partition of your hard drive at the top of the window, then mount it by clicking on the *Mount* button. Mandriva usually puts this at `/mnt/windows`.

Your fonts folder is probably `/mnt/windows/windows/fonts` if `/mnt/windows` is where your Windows partition is mounted.

Edit the file `/etc/X11/fs/config` (as the root user, of course), and append the following line:

```
/mnt/windows/windows/fonts
```

(or the appropriate directory, if yours was different). Be sure to put a comma at the end of the previous line. Save the file!

As the root user, execute the following command to restart the font server:

```
/etc/rc.d/init.d/xfs restart
```

You should be good to go: any TrueType fonts in your Windows partition will now be available to X Windows.

—Arun Sakthian M.R., arunmr87@gmail.com

Hardware specifications

Here's a tip that will give you the hardware details of your system. Run the following command as the root user:

```
lshw
```

We can get specific details by using the `-C` flag. For example:

```
lshw -C disk
```

...will list all your hard disks.

It can even create an HTML page with your hardware details if you use the following command:

```
lshw -html > My_Hardware_Spec.html
```

—Ravi ravi.shivanarayana@gmail.com

List all IPs

If you have a lot of network interfaces configured on a system and you want to view the IP addresses in a short list on it, execute the following command:

```
ip addr list
```

—Trupti Raul, truptii.raul@gmail.com

Run a command repeatedly and display the output

`watch` runs a command repeatedly, displaying its output (the first screen-full). This allows you to watch the program output change over time. By default, the program is run every two seconds. `watch` is very similar to `tail`.

```
watch -d ls -l
```

—Amey Parulekar, kindled.blade@gmail.com

Back up your HDD to another host

You can back up your hard disk to another host via `ssh` as the root, using the following command:

```
dd bs=1M if=/dev/sda | gzip | ssh \
username@remote-ip-address 'dd of=hda.gz'
```

—Ajeet Singh Raina, ajeet.singh.raina@logica.com

Extracting a file name from a path name

If you have a variable `$FILE` that contains a full path and file name, for example, `/home/fred/myfile.txt`, then the following code:

```
${FILE##*/}
```

...will evaluate to the file name without the path, for example, `myfile.txt`. The `##` string operator deletes the longest match of the specified regular expression.

—Ajeet Singh Raina, ajeet.singh.raina@logica.com

New look to your boot loader

If you are a little creative, you can have the bootloader look the way you want it to by changing the Splash image. However, images in JPEG or other popular formats won't do. For this, you will have to convert it into `.xpm` format (GIMP's default). Open a terminal and enter the following command:

```
convert - depth 8 -colors 14 -resize 640x480 \  
/root/grub-pics/photo.jpg /root/grub-pics/photo.xpm
```

Of course, change `/root/grub-pics/photo.xpm` to the path where your image is stored. Now compress it to the `gz` format by running the following code:

```
cd /root/grub-pics  
gzip photo.xpm
```

The output of the command is the file we need: `photo.xpm.gz`. Now, copy the image to your `/boot/grub` directory:

```
cp /root/grub-pics/photo.xpm.gz /boot/grub
```

The final step is to edit the Grub config file. Open the `grub.conf` (in Fedora, or `menu.lst` in most other systems) in the `/boot/grub` directory using a text editor and for the following line:

```
splashimage=
```

...change the path of the Splash screen image from `/grub/splash.xpm.gz` to:

```
/grub/photo.xpm.gz
```

Reboot the system and there you have your own customised boot Splash screen.

—Subramanian. M, supersubru@gmail.com

Create a quick Yum repository

The following are the steps to make a Yum repo:

1. Become the root user and edit the `/etc/yum.conf` file by changing the following line:

```
keepcache=0
```

```
...to
```

```
keepcache=1
```

2. Update packages, get new packages with `Yum` or `Pirut`. Switch back to being the normal user.
3. Make a directory called `packages` in your home directory. Go to `/var/cache/yum` and copy all the updated rpm packages from the sub-folders into `packages`.
4. Run the following command:

```
createrepo -pd ~/packages
```

A directory called `repodata` will be created inside the `packages` directory.

5. Burn the packages directory in the root of a CD or DVD. Name the CD/DVD `my_repo`.
6. Make a text file `local.repo` with the text below:

```
[local]  
name = My Repository  
baseurl = file:///media/my_repo/packages  
gpgcheck = 0  
enabled = 1
```

...where, `/media/my_repo` is the mount point of the media. Change if required.

7. Copy the `local.repo` file into `/etc/yum.repos.d` directory.

Now you can access your local media repository as usual.

Note: Change the `baseurl` appropriately, so that it points to the directory containing `repodata`, to keep the repository in the HDD. Change the name field accordingly. Use a CD/DVD RW, so that you can update the medium. The filename and the file heading in the `.repo` file should be the same.

—Arjun Pakrashi, phoxis@rediffmail.com



Share Your Linux Recipes!

The joy of using Linux is in finding ways to get around problems—take them head on, defeat them! We invite you to share your tips and tricks with us for publication in LFY so that they can reach a wider audience. Your tips could be related to administration, programming, troubleshooting or general tweaking. Submit them at <http://www.linuxforu.com>

The sender of each published tip will get an LFY T-shirt.

FOSS Yellow Pages


The best place for you to buy and sell FOSS products and services

HIGHLIGHTS

- A cost-effective marketing tool
- A user-friendly format for customers to contact you
- A dedicated section with yellow back-ground, and hence will stand out
- Reaches to tech-savvy IT implementers and software developers
- 80% of LFY readers are either decision influencers or decision takers
- Discounts for listing under multiple categories
- Discounts for booking multiple issues

FEATURES

- Listing is categorised on the basis of products and services
- Complete contact details plus 30-word description of organisation
- Option to print the LOGO of the organisation too (extra cost)
- Option to change the organisation description for listings under different categories



TARIFF

Category Listing

ONE Category	Rs 2,000
TWO Categories.....	Rs 3,500
THREE Categories.....	Rs 4,750
ADDITIONAL Category	Rs 1,000

Value-add Options


LOGO-plus-Entry.....	Rs 500
Highlight Entry (white background).....	Rs 1,000
Per EXTRA word (beyond 30 words).....	Rs 50

KEY POINTS

- Above rates are per-category basis.
- Above rates are charges for publishing in a single issue of LFY.
- Max. No. of Words for Organisation Description: 30 words.


TERMS & CONDITIONS

- Fill the form (below).
- You can use multiple copies of the form for multiple listings under different categories.
- Payment to be received along with booking.



Tear & Send

ORDER FORM



Tear & Send

Organisation Name (70 characters): _____

Description (30 words): _____

Email: _____ Website: _____

STD Code: _____ Phone: _____ Mobile: _____

Address (will not be published): _____

City/Town: _____ Pin-code: _____

CATEGORIES

- | | | |
|---|--|---|
| <input type="checkbox"/> CONSULTANTS | <input type="checkbox"/> HIGH PERFORMANCE COMPUTING | <input type="checkbox"/> SOFTWARE DEVELOPMENT |
| <input type="checkbox"/> CONSULTANT (FIRM) | <input type="checkbox"/> IT INFRASTRUCTURE SOLUTIONS | <input type="checkbox"/> TRAINING FOR PROFESSIONALS |
| <input type="checkbox"/> EMBEDDED SOLUTIONS | <input type="checkbox"/> LINUX-BASED WEB-HOSTING | <input type="checkbox"/> TRAINING FOR CORPORATE |
| <input type="checkbox"/> ENTERPRISE COMMUNICATION SOLUTIONS | <input type="checkbox"/> MOBILE SOLUTIONS | <input type="checkbox"/> THIN CLIENT SOLUTIONS |

Please find enclosed a sum of Rs. _____ by DD/ MO//crossed cheque* bearing the No. _____ dt. _____ in favour of


EFY Enterprises Pvt Ltd, payable at Delhi. (*Please add Rs. 50 on non-metro cheque) towards the cost of _____ FOSS Yellow Pages advertisement(s)

or charge my credit card ☐ VISA ☐ Master Card Please charge Rs. _____

against my credit card No. _____

C V V No. _____ (Mandatory)

Date of Birth ____ / ____ / ____ (dd/mm/yy) Card Expiry Date ____ / ____ (mm/yy)



EFY Enterprises Pvt Ltd., D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 020
Ph: 011-26810601-03, Fax: 011-26817565, Email: info@efyindia.com; Website: www.efyindia.com

Signature (as on the card)

To Book Your Listing, Call: Somaiah (B'lore: 09986075717) Suraj (Delhi: 09350475845)

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Somaiah (Bangalore) **09986075717**, Abhay (Delhi) **09350256557**

Enterprise Communication Solutions



Netcore Solutions Pvt. Ltd.

Emergic CleanMail is an Internet-based filtering service to monitor and filter e-mail traffic to protect against virus attacks, spam mails and wasted bandwidth. It offers triple-level virus scanning with extensive graphical reports on mail usage patterns. Emergic MailArchiva - Mail Archiving solutions for corporates which can get integrated with all mailing platform.

Mumbai-400013

Phone: 022-66628000

Mobile: 09322985222

Email: kalpit@netcore.co.in

Web: www.netcore.co.in



Red Hat India Pvt. Ltd.

Red Hat is the world's leading open source solutions provider. Red Hat provides high-quality, affordable technology with its operating system platform, Red Hat Enterprise Linux, together with applications, management and Services Oriented Architecture (SOA) solutions, including JBoss Enterprise Middleware. Red Hat also offers support, training and consulting services to its customers worldwide.

Mumbai-400076

Phone: 022-39878888

Email: marketing-in@redhat.com

Web: www.redhat.in

IT Infrastructure Solutions

BakBone Software Inc.

BakBone Software Inc. delivers complexity-reducing data protection technologies, including award-winning Linux solutions; proven Solaris products; and application-focused Windows offerings that reliably protect MS SQL, Oracle, Exchange, MySQL and other business critical applications.

New Delhi-110048

Tel: 011-42235156

Email: ashish.gupta@bakbone.com

Web: www.bakbone.com

Keen & Able Computers Pvt. Ltd.

Open Source Solutions Provider. Red Hat Ready Business Partner. Mail Servers/Anti-spam/GUI interface/Encryption, Clustering & Load Balancing - SAP/Oracle/Web/Thin Clients, Network and Host Monitoring, Security Consulting, Solutions, Staffing and Support.

New Delhi-110019

Tel: 011-30880046, 30880047

Mobile: 09810477448, 09891074905

Email: info@keenable.com

Web: www.keenable.com



Red Hat India Pvt. Ltd.

Red Hat is the world's leading open source solutions provider. Red Hat provides high-quality, affordable technology with its operating system platform, Red Hat Enterprise Linux, together with applications, management and Services Oriented Architecture (SOA) solutions, including JBoss Enterprise Middleware. Red Hat also offers support, training and consulting services to its customers worldwide.

Mumbai-400076

Phone: 022-39878888

Email: marketing-in@redhat.com

Web: www.redhat.in

TechnoInfotech™

A company focussed on Enterprise Solution using opensource software. Key Solutions:

- Enterprise Email Solution
- Internet Security and Access Control
- Managed Services for Email Infrastructure.

Mumbai-400001

Phone: 022-66338900; Extn. 324

Email: sales@technoinfotech.com

Web: www.technoinfotech.com

Software Development

Carizen Software (P) Ltd.

Carizen's flagship product is Rainmail Intranet Server, a complete integrated software product consisting modules like mail sever, proxy server, gateway anti-virus scanner, anti-spam, groupware, bandwidth aggregator & manager, firewall, chat server and fax server. Infrastructure.

Chennai-600028

Phone: 044-24958222, 8228, 9296

Email: info@carizen.com

Web: www.carizen.com

Thin Client Solutions

Enjay Network Solutions

Gujarat based ThinClient Solution Provider. Providing Small Size ThinClient PCs & a Full Featured ThinClient OS to perfectly suite needs of different working environment. Active Dealer Channel all over India.

Gujrat-396105

Tel.: 0260-3203400, 3241732, 3251732,

Mobile: 09377107650, 09898007650

Email: info@enjayworld.com

Web: www.enjayworld.com

Training for Corporate

Complete Open Source Solutions

RHCT, RHCE and RHCSS training. Hyderabad-500038

Phone: 040-66773365, 9849742065

Email: nayak.sujeet@gmail.com

Web: www.cossindia.com



FOSTERing Linux

Linux & Open Source Training Institute. All trainings provided by experienced experts & System Administrators only, RHCE, RHCSS, (Red Hat Training & Examination Partners), PHP, Perl, OpenOffice, Clustering, Mail Servers, Bridging the GAP by providing: Quality training (corporate & individual), Quality Manpower, Staffing and Support & 100% Placement Assistance.

Gurgaon-122001

Tel: 0124-4268187, 4080880

Mobile: 09350640169, 09818478555

Email: info@fosteringlinux.com

Web: www.fl.keenable.com

Lynus Academy Pvt. Ltd.

India's premier Linux and OSS training institute.

Chennai-600101

Phone: 044-42171278, 9840880558

Email: contactus@lynusacademy.com

Web: www.lynusacademy.com

To advertise in this section, please contact

Somaiah (Bangalore) **09986075717**

Abhay (Delhi) **09350256557**

Linux Learning Centre Private Limited

Pioneers in training on Linux technologies.

Bangalore-560019

Phone: 080-22428538, 26600839

Email: info@linuxlearningcentre.com

Web: www.linuxlearningcentre.com

Netweb Technologies

Simplified and scalable storage solutions.

Bangalore-560001

Phone: 080-41146565, 32719516

Email: info@netwebindia.com

Web: www.netwebindia.com

New Horizons India Ltd.

New Horizons India Ltd., a joint venture of New Horizons Worldwide, Inc. (NASDAQ: NEWH) and the Shriram group, is an Indian company operational since 2002 with a global foot print engaged in the business of knowledge delivery through acquiring, creating, developing, managing, lending and licensing knowledge in the areas of IT, Applied Learning. Technology Services and Supplementary Education. The company has pan India presence with 15 offices and employs 750 people.

New Delhi-110003

Tel: 011-43612400

Email: info@nhindia.com

Web: www.nhindia.com

Training for Professional

IPSR Solutions Ltd

Earn RHCE/RHCSS certification in Kerala along with a boating u free accomodation. IPSR conducted more than 2000 RHCE exams with 95-100% pass rate. Our faculty panel consists of Red Hat certified engineers.

ML Road, Kottayam

Mobile: +91-9447294635

Email: training@ipsrsolutions.com

Web: www.ipsr.org

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

No one brews Coffee & Conversation Like Karan.



All the madness, all the masala, all the gossip and all the controversies now captured in special edition books. Relive those fun moments with your favourite stars where they revealed it all. Read through those candid interviews that provoked conflicts and kindled friendships in tinsel town.

Review on People Magazine on Aug 15, 2008



“Koffee with Karan is a racy and fun read. The book is well laid out with enticing blurbs that lure you to read on.”

Mr. Bachchan has been the biggest ever superstar in Hindi Cinema

-Shahrukh Khan

I washed everyone's clothes including underwears

-Rakhi Sawant

I go to the gym so that I can watch men working out!

-Lara Dutta

Size does matter!

-Malaika Arora

I was great in bed


-Saif Ali Khan



Available in
Hindi & English

Shop online at
www.dpb.in

STAR ★ Koffee with Karan Logo and all associated art © STAR India Pvt Ltd. All Rights Reserved.




Accelerate your IT Capabilities. Put Brakes on Cost.

Welcome to Microsoft Virtualization.

Microsoft Virtualization breaks down barriers to creating the Virtual Enterprise. With end-to-end solutions, not only can you manage your technology infrastructure in an easy, smart and flexible manner, but also accelerate IT capabilities while reducing costs.

Go Virtual. Now!



INTEROP[®]

MUMBAI | 4-6 FEBRUARY, 2009

MMRDA GROUNDS, BANDRA KURLA COMPLEX

THE LEADING BUSINESS TECHNOLOGY EVENT

LAS VEGAS • MUMBAI • NEW YORK • SÃO PAULO • TOKYO

CALL FOR
PAPERS
OPEN

Learn about all the technologies that drive your business:

- Data Center
- Enterprise 2.0
- Wireless and Mobility
- Green IT
- Software-as-a-Service
- Cloud Computing
- SOA
- Open Source
- Document Management
- Application Delivery
- Security
- Storage
- Unified Communications
- Virtualization
- IP Telephony
- Telepresence
- Networking and Services
- Business Intelligence



For exhibiting and sponsorship information contact
Pankaj Jain at pankajj@ubmindia.com or call
(Mobile) + 91 98101 72077 or (Tel) + 91 11 2376 5553

Registration will open in November. To join the mailing
list, visit www.interop.com/mumbai/list

THE LEADING GLOBAL BUSINESS TECHNOLOGY EVENT IS COMING TO INDIA

Interop drives the growth and innovation of the global business technology market. As IT's most comprehensive event, Interop brings together the largest audience of enterprise buyers, SMB buyers, channel buyers, press and analysts—all in one place—to learn about the latest technologies, products and services.

Business and IT professionals attend Interop to get the information they need to purchase the latest technology. Tomorrow's market leaders are defined at Interop. Be part of the excitement this February.



MEDIA PARTNERS:



SUPPORTING ASSOCIATION:



United Business Media

©2008 CMP Media LLC. All Rights Reserved.
A United Business Media company.